Reference guide

# AWS General Reference

**Version 1.0**

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

# AWS General Reference: Reference guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

# Table of Contents

**AWS General Reference ..... 1**
AWS security credentials ..... 1
AWS IP address ranges ..... 1
AWS APIs ..... 2
AWS services endpoints and quotas ..... 2
AWS Glossary ..... 2
**AWS service endpoints ..... 3**
Regional endpoints ..... 3
Global endpoints ..... 6
View the service endpoints ..... 6
FIPS endpoints ..... 7
Minimum TLS version for FIPS endpoints ..... 7
Dual stack endpoints ..... 7
Learn more ..... 8
**AWS service quotas ..... 9**
**Service endpoints and quotas ..... 11**
Account access management ..... 20
Service endpoints ..... 21
Service quotas ..... 23
Amazon AI Operations ..... 23
Service endpoints ..... 23
Service quotas ..... 25
Amplify ..... 25
Amplify endpoints ..... 25
Amplify Studio (backend) endpoints ..... 28
Amplify Studio (UI Builder) endpoints ..... 30
Amplify Service quotas ..... 32
Amplify Studio (UI Builder) Service quotas ..... 35
API Gateway ..... 36
Service endpoints ..... 36
Service quotas ..... 46
AWS AppConfig ..... 55
Service endpoints ..... 55
Service quotas ..... 66

AppFabric ... 71
Service endpoints ... 71
Service quotas ... 72
ARC ... 72
Service endpoints ... 73
Service quotas ... 84
App Mesh ... 84
Service endpoints ... 84
Service quotas ... 95
App Runner ... 96
Service endpoints ... 97
Service quotas ... 98
Amazon AppFlow ... 100
Service endpoints ... 100
Service quotas ... 102
Application Auto Scaling ... 110
Service endpoints ... 110
Service quotas ... 115
Application Discovery Service ... 119
Service endpoints ... 120
Service quotas ... 121
MGN ... 122
Service endpoints ... 122
Service quotas ... 127
WorkSpaces Applications ... 131
Service endpoints ... 131
Service quotas ... 135
AWS App Studio ... 277
Service endpoints ... 278
Service quotas ... 278
AWS AppSync ... 278
Service endpoints ... 278
Service quotas ... 289
Athena ... 296
Service endpoints ... 296
Service quotas ... 301

Audit Manager ..... 319
Service endpoints ..... 319
Service quotas ..... 321
Amazon A2I ..... 321
Service endpoints ..... 322
Service quotas ..... 323
Aurora ..... 324
Service endpoints ..... 324
Service quotas ..... 332
Aurora DSQL ..... 337
Service endpoints ..... 338
Service quotas ..... 340
AWS Auto Scaling ..... 341
Service endpoints ..... 341
Service quotas ..... 344
AWS MCP Server ..... 345
Service endpoints ..... 345
Service quotas ..... 346
B2B Data Interchange ..... 346
Service endpoints ..... 346
Service quotas ..... 348
AWS Backup ..... 348
Service endpoints ..... 349
Service quotas ..... 353
AWS Batch ..... 355
Service endpoints ..... 356
Service quotas ..... 361
Amazon Bedrock ..... 363
Amazon Bedrock service endpoints ..... 363
Amazon Bedrock service quotas ..... 377
Amazon Bedrock AgentCore ..... 751
Service endpoints ..... 751
Service quotas ..... 755
Billing and Cost Management ..... 755
Service endpoints ..... 756
Service quotas ..... 762

AWS Billing Conductor ... 764
Service endpoints ... 765
Service quotas ... 765
Braket ... 765
Service endpoints ... 765
Service quotas ... 766
AWS Certificate Manager ... 787
Service endpoints ... 787
Service quotas ... 792
Amazon Chime ... 798
Service endpoints ... 799
Service quotas ... 799
Amazon Chime SDK ... 799
Service endpoints ... 800
Service quotas ... 808
AWS Clean Rooms ... 828
Service endpoints ... 828
Service quotas ... 830
AWS Clean Rooms ML ... 845
Service endpoints ... 845
Service quotas ... 847
Cloud Control API ... 872
Service endpoints ... 873
Service quotas ... 878
AWS Cloud9 ... 878
Service endpoints ... 879
Service quotas ... 882
Cloud Directory ... 883
Service endpoints ... 884
Service quotas ... 885
Cloud WAN ... 885
Service endpoints ... 885
Service quotas ... 893
CloudFormation ... 893
Service endpoints ... 894
Service quotas ... 898

CloudFront ... 917
Service endpoints ... 917
Service quotas ... 918
AWS CloudHSM ... 933
Service endpoints ... 933
Service quotas ... 938
AWS Cloud Map ... 939
Service endpoints ... 939
Service quotas ... 945
Amazon CloudSearch ... 946
Service endpoints ... 947
Service quotas ... 948
CloudShell ... 949
Service endpoints ... 950
Service quotas ... 953
Amazon CodeCatalyst ... 954
Service endpoints ... 955
Service quotas ... 955
AWS Control Catalog ... 955
Service Endpoints ... 955
Service Quotas ... 960
CloudTrail ... 961
Service endpoints ... 961
Service quotas ... 968
CloudWatch ... 980
Service endpoints ... 980
Service quotas ... 990
CloudWatch Application Insights ... 1002
Service endpoints ... 1002
Service quotas ... 1007
CloudWatch Application Signals ... 1008
Service quotas ... 1012
CloudWatch Events ... 1016
Service endpoints ... 1016
Service quotas ... 1022
CloudWatch Internet Monitor ... 1028

Service endpoints ..... 1029
Service quotas ..... 1033
CloudWatch Logs ..... 1033
Logs Service endpoints ..... 1033
Service quotas ..... 1039
CloudWatch Network Synthetic Monitor ..... 1050
Service endpoints ..... 1050
Service quotas ..... 1054
CloudWatch OAM ..... 1055
Service endpoints ..... 1055
Service quotas ..... 1060
CloudWatch Observability Admin ..... 1064
Service endpoints ..... 1064
Service quotas ..... 1068
CloudWatch RUM ..... 1074
Service endpoints ..... 1074
Service quotas ..... 1078
CloudWatch Synthetics ..... 1079
Service endpoints ..... 1079
Service quotas ..... 1084
CodeArtifact ..... 1084
Service endpoints ..... 1085
Service quotas ..... 1086
CodeBuild ..... 1088
Service endpoints ..... 1089
Service quotas ..... 1093
CodeCommit ..... 1099
Service endpoints ..... 1099
Service quotas ..... 1102
CodeConnections ..... 1103
Service endpoints ..... 1103
Service quotas ..... 1105
CodeDeploy ..... 1105
Service endpoints ..... 1106
CodeDeploy agent endpoints ..... 1110
Service quotas ..... 1121

CodeGuru Profiler ..... 1127
Service endpoints ..... 1127
Service quotas ..... 1129
CodeGuru Reviewer ..... 1129
Service endpoints ..... 1129
Service quotas ..... 1130
CodeGuru Security ..... 1131
Service endpoints ..... 1131
CodePipeline ..... 1132
Service endpoints ..... 1133
Service quotas ..... 1136
AWS CodeStar Connections ..... 1141
Service endpoints ..... 1141
Service quotas ..... 1143
AWS CodeStar Notifications ..... 1143
Service endpoints ..... 1144
Service quotas ..... 1146
Amazon Cognito ..... 1146
Service endpoints ..... 1146
Service quotas ..... 1156
Amazon Comprehend ..... 1172
Service endpoints ..... 1172
Service quotas ..... 1174
Amazon Comprehend Medical ..... 1190
Service endpoints ..... 1190
Service quotas ..... 1192
Compute Optimizer ..... 1198
Service endpoints ..... 1198
Service quotas ..... 1202
AWS Config ..... 1202
Service endpoints ..... 1203
Service quotas ..... 1207
Connect Customer ..... 1209
Service endpoints ..... 1209
Service quotas ..... 1220
Connect Health ..... 1381

Service endpoints ..... 1381
Service quotas ..... 1381
AWS Control Tower ..... 1381
Service Endpoints ..... 1381
Service Quotas ..... 1387
AWS Data Exchange ..... 1388
Service endpoints ..... 1388
Service quotas ..... 1389
Amazon DataZone ..... 1396
Service endpoints ..... 1396
Service quotas ..... 1398
Amazon Data Lifecycle Manager ..... 1442
Service endpoints ..... 1442
Service quotas ..... 1447
AWS Data Pipeline ..... 1447
Service endpoints ..... 1448
Service quotas ..... 1449
DataSync ..... 1450
Service endpoints ..... 1450
Service quotas ..... 1456
AWS DMS ..... 1459
Service endpoints ..... 1460
Service quotas ..... 1464
AWS Deadline Cloud ..... 1467
Service endpoints ..... 1468
Service quotas ..... 1471
AWS DeepRacer ..... 1477
Service endpoints ..... 1477
Service quotas ..... 1477
Detective ..... 1477
Service endpoints ..... 1478
Service quotas ..... 1481
DevOps Guru ..... 1481
Service endpoints ..... 1482
Service quotas ..... 1484
Device Farm ..... 1484

Service endpoints ..... 1485
Service quotas ..... 1485
AWS Direct Connect ..... 1486
Service endpoints ..... 1486
Service quotas ..... 1490
Directory Service ..... 1493
Directory Service endpoints ..... 1493
Directory Service Data endpoints ..... 1498
Service quotas ..... 1500
Amazon DocumentDB ..... 1501
Service endpoints ..... 1502
Service quotas ..... 1506
DynamoDB ..... 1507
Service endpoints ..... 1508
Service quotas ..... 1531
Elastic Beanstalk ..... 1537
Service endpoints ..... 1538
Service quotas ..... 1548
Amazon EBS ..... 1549
Service endpoints ..... 1549
Service quotas ..... 1555
Amazon EC2 ..... 1579
Service endpoints ..... 1579
Service quotas ..... 1585
Amazon EC2 Auto Scaling ..... 1588
Service endpoints ..... 1589
Service quotas ..... 1598
EC2 Image Builder ..... 1600
Service endpoints ..... 1600
Service quotas ..... 1605
EC2 Instance Connect ..... 1608
Service endpoints ..... 1609
Service quotas ..... 1613
Amazon ECR ..... 1613
Service endpoints ..... 1614
Service quotas ..... 1625

Amazon ECR Public ..... 1631
Service endpoints ..... 1631
Service quotas ..... 1631
Amazon ECS ..... 1636
Service endpoints ..... 1636
Service quotas ..... 1641
AWS Fargate quotas ..... 1652
Amazon EKS ..... 1659
Service endpoints ..... 1660
Service quotas ..... 1668
AWS Fargate service quotas ..... 1671
Amazon EFS ..... 1672
Service endpoints ..... 1672
Service quotas ..... 1681
Elastic Load Balancing ..... 1686
Service endpoints ..... 1687
Service quotas ..... 1692
Elastic Disaster Recovery ..... 1694
Service endpoints ..... 1695
Service quotas ..... 1700
ElastiCache ..... 1701
Service endpoints ..... 1701
Service quotas ..... 1706
Elemental Inference ..... 1708
Service endpoints ..... 1708
Service quotas ..... 1709
Amazon MemoryDB ..... 1713
Service endpoints ..... 1713
Service quotas ..... 1716
AWS Entity Resolution ..... 1716
Service endpoints ..... 1717
Service quotas ..... 1719
AWS End User Messaging ..... 1723
Service endpoints ..... 1723
Service quotas ..... 1736
Amazon EMR ..... 1741

Service endpoints ... 1742
Service quotas ... 1747
EventBridge ... 1773
Service endpoints ... 1773
Service quotas ... 1778
EventBridge Pipes ... 1784
Service endpoints ... 1785
Service quotas ... 1788
EventBridge Scheduler ... 1790
Service endpoints ... 1790
Service quotas ... 1795
EventBridge Schemas ... 1804
Service endpoints ... 1805
Service quotas ... 1808
Amazon EVS ... 1809
Service endpoints ... 1809
Service quotas ... 1813
FinSpace ... 1814
Service quotas ... 1814
AWS FIS ... 1820
Service endpoints ... 1821
Service quotas ... 1824
Firehose ... 1837
Service endpoints ... 1837
Service quotas ... 1842
Firewall Manager ... 1846
Service endpoints ... 1847
Service quotas ... 1851
Forecast ... 1857
Service endpoints ... 1858
Service quotas ... 1860
Amazon Fraud Detector ... 1866
Service endpoints ... 1867
Service quotas ... 1867
FreeRTOS ... 1870
Service endpoints ... 1870

Service quotas ..... 1874
Amazon FSx ..... 1874
Service endpoints ..... 1875
Service quotas ..... 1879
Amazon GameLift Servers ..... 1886
Service endpoints ..... 1887
UDP ping beacon endpoints ..... 1889
Service quotas ..... 1892
Amazon GameLift Streams ..... 1896
Service endpoints ..... 1897
Service quotas ..... 1897
Global Accelerator ..... 1914
Service endpoint (global) ..... 1914
Service quotas ..... 1915
Global Networks for Transit Gateways ..... 1917
Service endpoints ..... 1918
Service quotas ..... 1926
AWS Glue ..... 1927
Service endpoints ..... 1927
Service quotas ..... 1932
DataBrew ..... 1938
Service endpoints ..... 1939
Service quotas ..... 1941
Amazon Managed Grafana ..... 1943
Service endpoints ..... 1943
Service quotas ..... 1945
AWS Ground Station ..... 1949
Service endpoints ..... 1949
Service quotas ..... 1951
GuardDuty ..... 1952
Service endpoints ..... 1953
Service quotas ..... 1957
AWS Health ..... 1959
Service endpoints ..... 1959
Service quotas ..... 1960
HealthImaging ..... 1960

Service endpoints ... 1961
Service quotas ... 1963
HealthLake ... 1965
Service endpoints ... 1965
Service quotas ... 1966
AWS HealthOmics ... 1979
Service endpoints ... 1965
Quotas for AWS HealthOmics ... 1983
IAM ... 1987
Service endpoints ... 1988
Service quotas ... 1992
IAM Access Analyzer ... 1996
Service endpoints ... 1996
Service quotas ... 2001
IAM Identity Center ... 2003
Service endpoints ... 2004
Service quotas ... 2018
IAM Roles Anywhere ... 2019
Service endpoints ... 2020
Service quotas ... 2024
Incident Manager ... 2027
Service endpoints ... 2027
Service quotas ... 2034
Amazon Inspector ... 2045
Service endpoints ... 2046
Service quotas ... 2054
Amazon Inspector Classic ... 2054
Service endpoints ... 2055
Service quotas ... 2056
AWS IoT Core ... 2057
Service endpoints ... 2058
Service quotas ... 2069
AWS IoT Device Defender ... 2171
Service endpoints ... 2171
Service quotas ... 2175
AWS IoT Device Management ... 2190

Service endpoints ... 2190
Service quotas ... 2210
AWS IoT FleetWise ... 2245
Service endpoints ... 2245
Service quotas ... 2246
AWS IoT Wireless ... 2250
Service endpoints ... 2250
Service quotas ... 2255
AWS IoT Greengrass V1 ... 2278
Service endpoints ... 2279
Service quotas ... 2286
AWS IoT Greengrass V2 ... 2288
Service endpoints ... 2288
Service quotas ... 2297
AWS IoT SiteWise ... 2299
Service endpoints ... 2299
Service quotas ... 2306
AWS IoT TwinMaker ... 2332
Service endpoints ... 2332
Service quotas ... 2335
AWS IoT TwinMaker API throttling limits ... 2338
Amazon IVS ... 2338
Service endpoints ... 2339
Service quotas ... 2342
Amazon Kendra ... 2350
Service endpoints ... 2350
Service quotas ... 2353
Amazon Keyspaces ... 2357
Service endpoints ... 2358
Service quotas ... 2363
AWS KMS ... 2368
Service endpoints ... 2368
Service quotas ... 2376
Kinesis Data Streams ... 2397
Service endpoints ... 2398
Service quotas ... 2403

Kinesis Video Streams ..... 2404
Service endpoints ..... 2404
Service quotas ..... 2407
Lake Formation ..... 2428
Service endpoints ..... 2428
Service quotas ..... 2434
Lambda ..... 2434
Service endpoints ..... 2435
Service quotas ..... 2439
Launch Wizard ..... 2453
Service endpoints ..... 1600
Service quotas ..... 2458
Amazon Lex ..... 2458
V2 service endpoints ..... 2459
Service quotas ..... 2462
License Manager ..... 2464
Service endpoints ..... 2464
Service quotas ..... 2482
Lightsail ..... 2487
Service endpoints ..... 2488
Service quotas ..... 2490
Amazon Location Service ..... 2495
Service endpoints ..... 2495
Control plane actions ..... 2529
Service quotas ..... 2532
Lookout for Equipment ..... 2549
Service endpoints ..... 2549
Service quotas ..... 2549
Macie ..... 2554
Service endpoints ..... 2554
Service quotas ..... 2557
AWS Mainframe Modernization ..... 2564
Service endpoints ..... 2565
Service quotas ..... 2568
Amazon ML ..... 2568
Service endpoints ..... 2569

Service quotas ... 2569
Managed Blockchain ... 2571
Amazon Managed Blockchain (AMB) Access service endpoints and quotas ... 2572
Amazon Managed Blockchain (AMB) Query service endpoints and quotas ... 2574
Managed Service for Apache Flink ... 2575
Service endpoints ... 2576
Service quotas ... 2580
AWS Managed Services ... 2581
Service endpoints ... 2581
Service quotas ... 2581
AWS Management Console ... 2581
Service endpoints ... 2582
Service quotas ... 2586
Amazon MWAA ... 2586
Service endpoints ... 2587
Service quotas ... 2596
AWS Marketplace ... 2597
Service endpoints ... 2597
Service quotas ... 2603
Mechanical Turk ... 2603
Service endpoints ... 2603
Service quotas ... 2604
Amazon MSK ... 2604
Service endpoints ... 2604
Service quotas ... 2609
Amazon MSK Connect ... 2610
Service endpoints ... 2611
Service quotas ... 2616
MediaConnect ... 2616
Service endpoints ... 2617
Service quotas ... 2619
MediaConvert ... 2620
Service endpoints ... 2621
Service quotas ... 2623
MediaLive ... 2638
Service endpoints ... 2638

Service quotas ... 2641
MediaPackage ... 2645
Service endpoints ... 2646
Service quotas ... 2654
MediaTailor ... 2659
Service endpoints ... 2659
Service quotas ... 2662
Migration Hub ... 2666
Service endpoints ... 2666
Service quotas ... 2667
Migration Hub Orchestrator ... 2667
Service endpoints ... 2668
Service quotas ... 2668
AWS Migration Hub Refactor Spaces ... 2669
Service endpoints ... 2669
Service quotas ... 2671
Migration Hub Strategy Recommendations ... 2672
Service endpoints ... 2673
Service quotas ... 2674
Amazon Monitron ... 2674
Service endpoints ... 2674
Service quotas ... 2675
Amazon MQ ... 2676
Service endpoints ... 2676
Service quotas ... 2680
Neptune ... 2684
Service endpoints ... 2684
Service quotas ... 2689
Network Firewall ... 2690
Service endpoints ... 2691
Service quotas ... 2695
OpenSearch Service ... 2699
Service endpoints ... 2699
Service quotas ... 2711
Oracle Database@AWS ... 2720
Service endpoints ... 2720

Organizations ... 2720
Service endpoints ... 2721
Service quotas ... 2725
AWS Outposts ... 2730
Service endpoints ... 2730
Amazon S3 on Outposts ... 2734
Service quotas ... 2738
AWS Panorama ... 2738
Service endpoints ... 2739
Service quotas ... 2739
AWS Payment Cryptography ... 2743
Control plane endpoints ... 2743
Data plane endpoints ... 2746
Service quotas ... 2749
AWS PCS ... 2750
Service endpoints ... 2750
Service quotas ... 2753
Internal quotas ... 2754
Relevant quotas for other AWS services ... 2754
Amazon Personalize ... 2754
Service endpoints ... 2755
Service quotas ... 2761
Amazon Pinpoint ... 2773
Service endpoints ... 2774
Service quotas ... 2777
Amazon Polly ... 2799
Service endpoints ... 2799
Service quotas ... 2803
AWS Private CA ... 2811
AWS Private CA endpoints ... 2811
AWS Private CA Connector for Active Directory endpoints ... 2816
AWS Private CA Connector for SCEP endpoints ... 2820
AWS Private CA quotas ... 2825
AWS Private CA Connector for Active Directory quotas ... 2831
AWS Private CA Connector for SCEP quotas ... 2838
Amazon Managed Service for Prometheus ... 2841

Service endpoints ... 2842
Service quotas ... 2853
AWS Proton ... 2861
Service endpoints ... 2861
Service quotas ... 2863
Amazon Q Business ... 2864
Service regions and endpoints ... 2864
Service quotas ... 2865
Amazon Q Developer ... 2866
Amazon Q Developer Transform service quotas ... 2866
Amazon Q Developer Pro service quotas ... 2866
Quick Setup ... 2867
Service endpoints ... 2868
Service quotas ... 2870
Amazon Quick Sight ... 2873
Service endpoints ... 2874
Service quotas ... 2878
AWS RAM ... 2888
Service endpoints ... 2888
Service quotas ... 2893
Recycle Bin ... 2895
Service endpoints ... 2895
Service quotas ... 2900
Amazon Redshift ... 2901
Service endpoints ... 2901
Service quotas ... 2920
Amazon Rekognition ... 2921
Service endpoints ... 2921
Service quotas ... 2928
Amazon RDS ... 2945
Service endpoints ... 2946
Service quotas ... 2956
AWS re:Post Private ... 2961
Service endpoints ... 2961
Service quotas ... 2962
Resilience Hub ... 2966

Service endpoints ... 2966
Service quotas ... 2969
Resource Explorer ... 2974
Service endpoints ... 2975
Service quotas ... 2980
Resource Groups and Tagging ... 2981
AWS Resource Groups ... 2981
Tagging ... 2986
Red Hat OpenShift Service on AWS ... 2992
Service endpoints ... 2992
Service quotas ... 2994
Route 53 ... 2998
Service endpoints ... 2998
Service quotas ... 3013
SageMaker AI ... 3017
Service Endpoints ... 3017
Service Quotas ... 3033
Secrets Manager ... 3247
Service endpoints ... 3248
Service quotas ... 3252
Security Lake ... 3255
Service endpoints ... 3255
Security Hub CSPM ... 3258
Service endpoints ... 3259
Service quotas ... 3263
AWS Security Incident Response ... 3265
Service endpoints ... 3265
Service quotas ... 3268
Service Quotas ... 3272
Service endpoints ... 3273
Service quotas ... 3277
AWS Serverless Application Repository ... 3282
Service endpoints ... 3282
Service quotas ... 3285
AWS Service Catalog endpoints and quotas ... 3285
Service endpoints ... 3286

Service quotas ... 3289
Shield Advanced ... 3292
Service endpoints ... 3292
Service quotas ... 3297
Amazon S3 ... 3297
Service endpoints ... 3298
Service quotas ... 3332
Amazon Glacier ... 3336
Service endpoints ... 3337
Service quotas ... 3340
Amazon SES ... 3341
API endpoints ... 3341
SMTP endpoints ... 3346
Email Receiving endpoints ... 3351
DKIM domains ... 3353
Tracking domains ... 3354
Feedback endpoints ... 3356
SMTP Relay IP Ranges ... 3357
Service quotas ... 3359
AWS Signer ... 3360
Service endpoints supporting signing for AWS Lambda, Amazon ECR, and Amazon EKS ... 3360
Service endpoints supporting signing for FreeRTOS and AWS IoT Device Management ... 3370
Service quotas ... 3372
AWS Sign-In ... 3376
Service endpoints ... 3377
OAuth service endpoints ... 3381
Service quotas ... 3383
Amazon SimpleDB ... 3383
Service endpoints ... 3384
Service quotas ... 3385
Amazon SNS ... 3385
Service endpoints ... 3385
Service quotas ... 3394
Amazon SQS ... 3401
Service endpoints ... 3401
Service quotas ... 3413

AWS STS ... 3416
Service endpoints ... 3417
Service quotas ... 3422
Amazon SWF ... 3422
Service endpoints ... 3423
Service quotas ... 3428
Snow Family ... 3445
Service endpoints ... 3445
Service quotas ... 3451
Step Functions ... 3452
Service endpoints ... 3452
Service quotas ... 3461
Storage Gateway ... 3478
Service endpoints ... 3478
Service quotas ... 3485
Support ... 3487
Service endpoints ... 3488
Service quotas ... 3494
AWS Support App in Slack endpoints ... 3494
Systems Manager ... 3495
Service endpoints for Systems Manager ... 3495
Service quotas ... 3500
AWS Systems Manager for SAP ... 3520
Service endpoints ... 3521
Service quotas ... 3525
Amazon Textract ... 3525
Service endpoints ... 3526
Service quotas ... 3528
Adapters quotas ... 3530
Timestream ... 3531
Service endpoints ... 3531
Service quotas ... 3538
AWS TNB ... 3543
Service endpoints ... 3543
Service quotas ... 3544
AWS Transform ... 3545

AWS Transform service quotas ... 3545
Amazon Transcribe ... 3549
Service endpoints ... 3549
Service quotas ... 3558
Transfer Family ... 3575
Service endpoints ... 3576
Service quotas ... 3580
Amazon Translate ... 3585
Service endpoints ... 3585
Service quotas ... 3588
Trusted Advisor API ... 3589
Service endpoints ... 3589
User Notifications ... 3590
Service endpoints ... 3590
AWS User Notifications Contacts endpoints ... 3594
Service quotas ... 3594
AWS User Experience Customization ... 3595
Service endpoints ... 3596
Service quotas ... 3599
Verified Access ... 3600
Service endpoints ... 3600
Service quotas ... 3600
Verified Permissions ... 3601
Service endpoints ... 3601
Service quotas ... 3607
Amazon VPC ... 3612
Service endpoints ... 3612
Service quotas ... 3612
VPC Lattice ... 3619
Service endpoints ... 3620
Service quotas ... 3623
AWS WAF ... 3627
Service endpoints ... 3628
Service quotas ... 3632
AWS WAF Classic ... 3642
Service endpoints ... 3643

Service quotas ... 3647
AWS Well-Architected Tool ... 3650
Service endpoints ... 3650
Service quotas ... 3653
AWS Wickr ... 3655
Service endpoints ... 3655
Amazon WorkMail ... 3657
Service endpoints ... 3657
Email protocols endpoints ... 3658
Service quotas ... 3659
Amazon WorkSpaces ... 3659
Service endpoints ... 3660
Service quotas ... 3663
Amazon WorkSpaces Instances ... 3668
Service endpoints ... 3668
Service quotas ... 3670
Amazon WorkSpaces Secure Browser ... 3671
Service endpoints ... 3671
Service quotas ... 3673
X-Ray ... 3675
Service endpoints ... 3676
Service quotas ... 3680
**AWS Lifecycle changes ... 3682**
Services in Maintenance ... 3682
Services in Sunset ... 3687
Services in Full Shutdown ... 3690
**AWS Lifecycle documentation history ... 3694**

# AWS General Reference

The AWS General Reference provides AWS service endpoint and quota information for Amazon Web Services. Additionally, you can find links to other common topics.

**Contents**


- AWS security credentials
- AWS IP address ranges
- AWS APIs
- AWS services endpoints and quotas
- AWS Glossary


## AWS security credentials

When you interact with AWS, you specify your AWS *security credentials* to verify who you are and whether you have permission to access the resources that you are requesting. AWS uses the security credentials to authenticate and authorize your requests.

For more information, see the following resources:

- AWS security credentials in the *IAM User Guide*
- AWS security audit guidelines in the *IAM User Guide*

## AWS IP address ranges

AWS publishes its current IP address ranges in JSON format. You can download a `.json` file to view current ranges.

The IP address ranges that you bring to AWS through bring your own IP addresses (BYOIP) are not included in the `.json` file.

For more information, see the following resources:

- AWS IP address ranges in the *Amazon VPC User Guide*
- AWS services that support IPv6 in the *Amazon VPC User Guide*

# AWS APIs

The following pages provide information that is useful when using an AWS API:

- Retry behavior in the *AWS SDKs and Tools Reference Guide*
- Signing AWS API requests in the *IAM User Guide*

# AWS services endpoints and quotas

You can learn about the endpoints and service quotas in the following pages:

- AWS service endpoints
- AWS service quotas
- Service endpoints and quotas
- Enable or disable AWS Regions in your account in the *AWS Account Management Guide*

# AWS Glossary

For the latest AWS terminology, see the AWS Glossary.

# AWS service endpoints

To connect programmatically to an AWS service, you use an endpoint. An *endpoint* is the URL of the entry point for an AWS web service. The AWS SDKs and the AWS Command Line Interface (AWS CLI) automatically use the default endpoint for each service in an AWS Region. But you can specify an alternate endpoint for your API requests.

If a service supports Regions, the resources in each Region are independent of similar resources in other Regions. For example, you can create an Amazon EC2 instance or an Amazon SQS queue in one Region. When you do, the instance or queue is independent of instances or queues in all other Regions.

**Contents**


- Regional endpoints
- Global endpoints
- View the service endpoints
- FIPS endpoints
- Dual stack endpoints
- Learn more


## Regional endpoints

Most Amazon Web Services offer a Regional endpoint that you can use to make your requests. In general, these endpoints support IPv4 traffic and they use the following syntax.

```
protocol://service-code.region-code.amazonaws.com
```

For example, `https://dynamodb.us-west-2.amazonaws.com` is the endpoint for the Amazon DynamoDB service in the US West (Oregon) Region.

The following table lists the name and code of each Region.

| Name | Code |
| --- | --- |
| US East (Ohio) | `us-east-2` |

| Name | Code |
| --- | --- |
| US East (N. Virginia) | us-east-1 |
| US West (N. California) | us-west-1 |
| US West (Oregon) | us-west-2 |
| Africa (Cape Town) | af-south-1 |
| Asia Pacific (Hong Kong) | ap-east-1 |
| Asia Pacific (Hyderabad) | ap-south-2 |
| Asia Pacific (Jakarta) | ap-southeast-3 |
| Asia Pacific (Malaysia) | ap-southeast-5 |
| Asia Pacific (Melbourne) | ap-southeast-4 |
| Asia Pacific (Mumbai) | ap-south-1 |
| Asia Pacific (New Zealand) | ap-southeast-6 |
| Asia Pacific (Osaka) | ap-northeast-3 |
| Asia Pacific (Seoul) | ap-northeast-2 |
| Asia Pacific (Singapore) | ap-southeast-1 |
| Asia Pacific (Sydney) | ap-southeast-2 |
| Asia Pacific (Taipei) | ap-east-2 |
| Asia Pacific (Thailand) | ap-southeast-7 |
| Asia Pacific (Tokyo) | ap-northeast-1 |
| Canada (Central) | ca-central-1 |
| Canada West (Calgary) | ca-west-1 |

| Name | Code |
| --- | --- |
| Europe (Frankfurt) | eu-central-1 |
| Europe (Ireland) | eu-west-1 |
| Europe (London) | eu-west-2 |
| Europe (Milan) | eu-south-1 |
| Europe (Paris) | eu-west-3 |
| Europe (Spain) | eu-south-2 |
| Europe (Stockholm) | eu-north-1 |
| Europe (Zurich) | eu-central-2 |
| Israel (Tel Aviv) | il-central-1 |
| Mexico (Central) | mx-central-1 |
| Middle East (Bahrain) | me-south-1 |
| Middle East (UAE) | me-central-1 |
| South America (São Paulo) | sa-east-1 |
| AWS GovCloud (US-East) | us-gov-east-1 |
| AWS GovCloud (US-West) | us-gov-west-1 |

### General endpoints

The following services support Regional endpoints but also support a general endpoint that doesn't include a Region. When you use a general endpoint, AWS routes the API request to US East (N. Virginia) (`us-east-1`), which is the default Region for API calls.

- Amazon EC2 – ec2.amazonaws.com
- Amazon EC2 Auto Scaling – autoscaling.amazonaws.com

- Amazon EMR – elasticmapreduce.amazonaws.com

## Global endpoints

The following services each have a global endpoint that spans AWS Regions:

- AWS Cloud WAN
- Amazon CloudFront
- AWS Global Accelerator
- AWS Identity and Access Management (IAM)
- AWS Organizations
- Amazon Route 53
- AWS Shield Advanced
- AWS WAF Classic

Note that some of these services have additional endpoints, such as FIPS endpoints.

## View the service endpoints

You can view the AWS service endpoints using the following options:

- Open Service endpoints and quotas, search for the service name, and click the link to open the page for that service. To view the supported endpoints for all AWS services in the documentation without switching pages, view the information in the Service Endpoints and Quotas page in the PDF instead.
- To programmatically check for service availability using the SDK for Java, see Check for service availability in an Region in the *AWS SDK for Java Developer Guide*.
- To programmatically view Region and service information using Systems Manager, see Calling public parameters for AWS services, Regions, endpoints, and zones in parameter store in the *AWS Systems Manager User Guide*. For information about how to use public parameters, see Query for AWS Regions, Endpoints, and More Using AWS Systems Manager Parameter Store.
- To see the supported AWS services in each Region (without endpoints), see the Region Table.

# FIPS endpoints

Some AWS services offer endpoints that support Federal Information Processing Standard (FIPS) 140 in some Regions. Unlike standard AWS endpoints, FIPS endpoints use a TLS software library that complies with FIPS 140. These endpoints might be required by enterprises that interact with the United States government.

To specify a FIPS endpoint when you call an AWS operation, use a mechanism provided by the tool that you're using to make the call. For example, the AWS SDKs provide the following mechanisms to use FIPS endpoints:

- Set the `AWS_USE_FIPS_ENDPOINT` environment variable to `true`
- Add `use_fips_endpoint=true` to your `~/.aws/config` file

For information about FIPS endpoint support in AWS SDKs, see Dual-stack and FIPS endpoints in the *AWS SDKs and Tools Reference Guide*.

The AWS Command Line Interface supports the preceding mechanisms, and also provides the `--endpoint-url` option. The following example uses `--endpoint-url` to specify the FIPS endpoint for AWS Key Management Service (AWS KMS) in the US West (Oregon) Region.

```
aws kms create-key --endpoint-url https://kms-fips.us-west-2.amazonaws.com
```

For a list of FIPS endpoints, see FIPS endpoints by Service.

## Minimum TLS version for FIPS endpoints

With FIPS endpoints, the minimum requirement is TLS 1.2. We recommend TLS 1.3. To determine whether your applications were affected by this change, see this AWS Security Blog post.

# Dual stack endpoints

Some AWS services offer dual stack endpoints, so that you can access them using either IPv4 or IPv6 requests. In general, the syntax of a dual stack endpoint is as follows:

```
protocol://service-code.region-code.api.aws
```

However, Amazon S3 uses the following syntax for Amazon S3 dual stack endpoints:

```
protocol://service-code.dualstack.region-code.amazonaws.com
```

To make a request to a dual stack endpoint, use the mechanism provided by the tool or AWS SDK to specify the endpoint. For information about dual stack endpoint support in AWS SDKs, see Dual-stack and FIPS endpoints in the *AWS SDKs and Tools Reference Guide*.

The AWS CLI provides the `--endpoint-url` option. The following example uses `--endpoint-url` to specify the dual stack endpoint for Amazon EC2 in the US West (Oregon) Region.

```
aws ec2 describe-regions --region us-west-2 --endpoint-url https://ec2.us-
west-2.api.aws
```

To use the AWS CLI to configure dual-stack endpoints for all AWS services, see Set to use dual-stack endpoints for all AWS services in the *AWS Command Line Interface User Guide*.

For a list of services that support dual stack endpoints, see AWS services that support IPv6.

## Learn more

You can find endpoint information from the following sources:

- To enable Regions that are disabled by default, see Enable or disable AWS Regions in your account in the *AWS Account Management Reference Guide*.
- For information about the AWS services and endpoints available in the China Regions, see China (Beijing) Region Endpoints and China (Ningxia) Region Endpoints.

# AWS service quotas

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, but not all quotas can be increased.

**To view service quotas**

You can view service quotas by using the following options:

- From the documentation: Open the Service endpoints and quotas page in the documentation, search for the service name, and then click the link to go to the page for that service. To view the service quotas for all AWS services in the documentation without switching pages, view the information in the PDF Service endpoints and quotas page.
- From the console: Open the Service Quotas console. In the navigation pane, choose **AWS services**, and then select a service. For more information, see Viewing service quotas in the *Service Quotas User Guide*.
- From the AWS CLI: Use the list-service-quotas and list-aws-default-service-quotas AWS CLI commands. For instructions, see Viewing service quotas in the *Service Quotas User Guide*.

Your account's actual quota value may be less than the AWS default quota value if the account was recently created or if you use the account minimally.

**To request a quota increase**

Support might approve, deny, or partially approve your quota increase requests. Increases are not granted immediately. It might take a couple of days for your increase to take effect.

You can request a quota increase by using one of following options:

- From the console: Open the Service Quotas console. In the navigation pane, choose **AWS services**. Select a service, select a quota, and follow the directions to request a quota increase. For instructions, see Requesting a quota increase in the *Service Quotas User Guide*.
- From the AWS CLI: Use the request-service-quota-increase AWS CLI command. For instructions, see Requesting a quota increase in the *Service Quotas User Guide*.
- From a support case: If a service is not yet available in Service Quotas, use the AWS Support Center Console to create a service quota increase case. If the service is available in Service

Quotas, we highly recommend that you use the Service Quotas console instead of creating a support case.

# Service endpoints and quotas

The following pages describe the service endpoints and service quotas for AWS services. To connect programmatically to an AWS service, you use an endpoint. For more information, see AWS service endpoints. Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

Choose one of the following links to go to the page for that service. To view the service quotas for all AWS services in the documentation without switching pages, view the information in the Service endpoints and quotas page in the PDF instead.

**Services**

- Account access management endpoints and quotas
- Amazon AI Operations endpoints and quotas
- AWS Amplify endpoints and quotas
- Amazon API Gateway endpoints and quotas
- AWS AppConfig endpoints and quotas
- AWS AppFabric endpoints and quotas
- Amazon Application Recovery Controller (ARC) endpoints and quotas
- AWS App Mesh endpoints and quotas
- AWS App Runner endpoints and quotas
- Amazon AppFlow endpoints and quotas
- Application Auto Scaling endpoints and quotas
- AWS Application Discovery Service endpoints and quotas
- AWS Transform MGN endpoints and quotas
- Amazon AppStream 2.0 endpoints and quotas
- AWS App Studio endpoints and quotas
- AWS AppSync endpoints and quotas
- Amazon Athena endpoints and quotas
- AWS Audit Manager endpoints and quotas
- Amazon Augmented AI endpoints and quotas

- Amazon Aurora endpoints and quotas
- Amazon Aurora DSQL endpoints and quotas
- AWS Auto Scaling endpoints and quotas
- AWS MCP Server endpoints and quotas
- AWS B2B Data Interchange endpoints and quotas
- AWS Backup endpoints and quotas
- AWS Batch endpoints and quotas
- Amazon Bedrock endpoints and quotas
- Amazon Bedrock AgentCore endpoints and quotas
- AWS Billing and Cost Management endpoints and quotas
- AWS Billing Conductor endpoints and quotas
- Amazon Braket endpoints and quotas
- AWS Certificate Manager endpoints and quotas
- Amazon Chime endpoints and quotas
- Amazon Chime SDK endpoints and quotas
- AWS Clean Rooms endpoints and quotas
- AWS Clean Rooms ML endpoints and quotas
- Cloud Control API endpoints and quotas
- AWS Cloud9 endpoints and quotas
- Amazon Cloud Directory endpoints and quotas
- AWS Cloud WAN endpoints and quotas
- AWS CloudFormation endpoints and quotas
- Amazon CloudFront endpoints and quotas
- AWS CloudHSM endpoints and quotas
- AWS Cloud Map endpoints and quotas
- Amazon CloudSearch endpoints and quotas
- AWS CloudShell endpoints and quotas
- Amazon CodeCatalyst endpoints and quotas

- AWS Control Catalog endpoints and quotas
- AWS CloudTrail endpoints and quotas
- Amazon CloudWatch endpoints and quotas
- Amazon CloudWatch Application Insights endpoints and quotas
- Amazon CloudWatch Application Signals endpoints and quotas
- Amazon CloudWatch Events endpoints and quotas
- Internet Monitor endpoints and quotas
- Amazon CloudWatch Logs endpoints and quotas
- Network Synthetic Monitor
- CloudWatch Observability Access Manager endpoints and quotas
- Amazon CloudWatch Observability Admin endpoints and quotas
- Amazon CloudWatch RUM endpoints and quotas
- Amazon CloudWatch Synthetics endpoints and quotas
- AWS CodeArtifact endpoints and quotas
- AWS CodeBuild endpoints and quotas
- AWS CodeCommit endpoints and quotas
- CodeConnections endpoints and quotas
- AWS CodeDeploy endpoints and quotas
- Amazon CodeGuru Profiler endpoints and quotas
- Amazon CodeGuru Reviewer endpoints and quotas
- Amazon CodeGuru Security endpoints and quotas
- AWS CodePipeline endpoints and quotas
- AWS CodeStar Connections endpoints and quotas
- AWS CodeStar Notifications endpoints and quotas
- Amazon Cognito endpoints and quotas
- Amazon Comprehend endpoints and quotas
- Amazon Comprehend Medical endpoints and quotas
- AWS Compute Optimizer endpoints and quotas

- AWS Config endpoints and quotas
- Amazon Connect Customer endpoints and quotas
- Amazon Connect Health endpoints and quotas
- AWS Control Tower endpoints and quotas
- AWS Data Exchange endpoints and quotas
- Amazon DataZone endpoints and quotas
- Amazon Data Lifecycle Manager endpoints and quotas
- AWS Data Pipeline endpoints and quotas
- AWS DataSync endpoints and quotas
- AWS Database Migration Service endpoints and quotas
- AWS Deadline Cloud endpoints and quotas
- AWS DeepRacer endpoints and quotas
- Amazon Detective endpoints and quotas
- Amazon DevOps Guru endpoints and quotas
- AWS Device Farm endpoints and quotas
- AWS Direct Connect endpoints and quotas
- Directory Service endpoints and quotas
- Amazon DocumentDB endpoints and quotas
- Amazon DynamoDB endpoints and quotas
- AWS Elastic Beanstalk endpoints and quotas
- Amazon Elastic Block Store endpoints and quotas
- Amazon EC2 endpoints and quotas
- Amazon EC2 Auto Scaling endpoints and quotas
- EC2 Image Builder endpoints and quotas
- Amazon EC2 Instance Connect endpoints and quotas
- Amazon ECR endpoints and quotas
- Amazon ECR Public endpoints and quotas
- Amazon ECS endpoints and quotas

- Amazon Elastic Kubernetes Service endpoints and quotas
- Amazon Elastic File System endpoints and quotas
- Elastic Load Balancing endpoints and quotas
- AWS Elastic Disaster Recovery endpoints and quotas
- Amazon ElastiCache endpoints and quotas
- AWS Elemental Inference endpoints and quotas
- Amazon MemoryDB endpoints and quotas
- AWS Entity Resolution endpoints and quotas
- AWS End User Messaging endpoints and quotas
- Amazon EMR endpoints and quotas
- Amazon EventBridge endpoints and quotas
- Amazon EventBridge Pipes endpoints and quotas
- Amazon EventBridge Scheduler endpoints and quotas
- Amazon EventBridge Schemas endpoints and quotas
- Amazon Elastic VMware Service endpoints and quotas
- Amazon FinSpace quotas
- AWS Fault Injection Service endpoints and quotas
- Amazon Data Firehose endpoints and quotas
- AWS Firewall Manager endpoints and quotas
- Amazon Forecast endpoints and quotas
- Amazon Fraud Detector endpoints and quotas
- FreeRTOS endpoints and quotas
- Amazon FSx endpoints and quotas
- Amazon GameLift Servers endpoints and quotas
- Amazon GameLift Streams endpoints and quotas
- AWS Global Accelerator endpoints and quotas
- AWS Global Networks for Transit Gateways Quotas (Network Manager)
- AWS Glue endpoints and quotas
- AWS Glue DataBrew endpoints and quotas

- Amazon Managed Grafana endpoints and quotas
- AWS Ground Station endpoints and quotas
- Amazon GuardDuty endpoints and quotas
- AWS Health endpoints and quotas
- AWS HealthImaging endpoints and quotas
- AWS HealthLake endpoints and quotas
- Service quotas and endpoints for AWS HealthOmics
- AWS Identity and Access Management endpoints and quotas
- IAM Access Analyzer endpoints and quotas
- AWS IAM Identity Center endpoints and quotas
- AWS Identity and Access Management Roles Anywhere endpoints and quotas
- AWS Systems Manager Incident Manager endpoints and quotas
- Amazon Inspector endpoints and quotas
- Amazon Inspector Classic endpoints and quotas
- AWS IoT Core endpoints and quotas
- AWS IoT Device Defender endpoints and quotas
- AWS IoT Device Management endpoints and quotas
- AWS IoT FleetWise endpoints and quotas
- AWS IoT Wireless endpoints and quotas
- AWS IoT Greengrass V1 endpoints and quotas
- AWS IoT Greengrass V2 endpoints and quotas
- AWS IoT SiteWise endpoints and quotas
- AWS IoT TwinMaker endpoints and quotas
- Amazon Interactive Video Service endpoints and quotas
- Amazon Kendra endpoints and quotas
- Amazon Keyspaces (for Apache Cassandra) endpoints and quotas
- AWS Key Management Service endpoints and quotas
- Amazon Kinesis Data Streams endpoints and quotas

- Amazon Kinesis Video Streams endpoints and quotas
- AWS Lake Formation endpoints and quotas
- AWS Lambda endpoints and quotas
- AWS Launch Wizard endpoints and quotas
- Amazon Lex endpoints and quotas
- AWS License Manager endpoints and quotas
- Amazon Lightsail endpoints and quotas
- Amazon Location Service endpoints and quotas
- Amazon Lookout for Equipment endpoints and quotas
- Amazon Macie endpoints and quotas
- AWS Mainframe Modernization endpoints and quotas
- Amazon Machine Learning endpoints and quotas
- Amazon Managed Blockchain endpoints and quotas
- Amazon Managed Service for Apache Flink endpoints and quotas
- AWS Managed Services endpoints and quotas
- AWS Management Console endpoints and quotas
- Amazon Managed Workflows for Apache Airflow endpoints and quotas
- AWS Marketplace endpoints and quotas
- Amazon Mechanical Turk endpoints and quotas
- Amazon Managed Streaming for Apache Kafka endpoints and quotas
- Amazon MSK Connect endpoints and quotas
- AWS Elemental MediaConnect endpoints and quotas
- AWS Elemental MediaConvert endpoints and quotas
- AWS Elemental MediaLive endpoints and quotas
- AWS Elemental MediaPackage endpoints and quotas
- AWS Elemental MediaTailor endpoints and quotas
- AWS Migration Hub endpoints and quotas
- Migration Hub Orchestrator endpoints and quotas

- AWS Migration Hub Refactor Spaces endpoints and quotas
- Migration Hub Strategy Recommendations endpoints and quotas
- Amazon Monitron endpoints and quotas
- Amazon MQ endpoints and quotas
- Amazon Neptune endpoints and quotas
- AWS Network Firewall endpoints and quotas
- Amazon OpenSearch Service endpoints and quotas
- Oracle Database@AWS endpoints and quotas
- AWS Organizations endpoints and quotas
- AWS Outposts endpoints and quotas
- AWS Panorama endpoints and quotas
- AWS Payment Cryptography endpoints and quotas
- AWS Parallel Computing Service endpoints and quotas
- Amazon Personalize endpoints and quotas
- Amazon Pinpoint endpoints and quotas
- Amazon Polly endpoints and quotas
- AWS Private Certificate Authority endpoints and quotas
- Amazon Managed Service for Prometheus endpoints and quotas
- AWS Proton endpoints and quotas
- Amazon Q Business endpoints and quotas
- Amazon Q Developer endpoints and quotas
- Systems Manager Quick Setup endpoints and quotas
- Amazon Quick Sight
- AWS Resource Access Manager endpoints and quotas
- Recycle Bin endpoints and quotas
- Amazon Redshift endpoints and quotas
- Amazon Rekognition endpoints and quotas
- Amazon Relational Database Service endpoints and quotas
- AWS re:Post Private endpoints and quotas

- AWS Resilience Hub endpoints and quotas
- AWS Resource Explorer
- AWS Resource Groups and Tagging endpoints and quotas
- Red Hat OpenShift Service on AWS endpoints and quotas
- Amazon Route 53 endpoints and quotas
- Amazon SageMaker AI endpoints and quotas
- AWS Secrets Manager endpoints and quotas
- Amazon Security Lake endpoints
- AWS Security Hub CSPM endpoints and quotas
- AWS Security Incident Response endpoints and quotas
- Service Quotas endpoints and quotas
- AWS Serverless Application Repository endpoints and quotas
- AWS Service Catalog endpoints and quotas
- AWS Shield Advanced endpoints and quotas
- Amazon Simple Storage Service endpoints and quotas
- Amazon Glacier endpoints and quotas
- Amazon Simple Email Service endpoints and quotas
- AWS Signer endpoints and quotas
- AWS Sign-In endpoints and quotas
- Amazon SimpleDB endpoints and quotas
- Amazon Simple Notification Service endpoints and quotas
- Amazon Simple Queue Service endpoints and quotas
- AWS Security Token Service endpoints and quotas
- Amazon Simple Workflow Service endpoints and quotas
- AWS Snow Family endpoints and quotas
- AWS Step Functions endpoints and quotas
- AWS Storage Gateway endpoints and quotas
- AWS Support endpoints and quotas
- AWS Systems Manager endpoints and quotas

- AWS Systems Manager for SAP endpoints and quotas
- Amazon Textract endpoints and quotas
- Amazon Timestream endpoints and quotas
- AWS Telco Network Builder endpoints and quotas
- AWS Transform endpoints and quotas
- Amazon Transcribe endpoints and quotas
- AWS Transfer Family endpoints and quotas
- Amazon Translate endpoints and quotas
- AWS Trusted Advisor API endpoints and quotas
- AWS User Notifications endpoints and quotas
- AWS User Experience Customization endpoints and quotas
- AWS Verified Access endpoints and quotas
- Amazon Verified Permissions endpoints and quotas
- Amazon Virtual Private Cloud endpoints and quotas
- Amazon VPC Lattice endpoints and quotas
- AWS WAF endpoints and quotas
- AWS WAF Classic endpoints and quotas
- AWS Well-Architected Tool endpoints and quotas
- AWS Wickr endpoints and quotas
- Amazon WorkMail endpoints and quotas
- Amazon WorkSpaces endpoints and quotas
- Amazon WorkSpaces Instances endpoints and quotas
- Amazon WorkSpaces Secure Browser endpoints and quotas
- AWS X-Ray endpoints and quotas

# Account access management endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | account-access.us-east-2.api.aws | | |
| US East (N. Virginia) | us-east-1 | account-access.us-east-1.api.aws | | |
| US West (N. California) | us-west-1 | account-access.us-west-1.api.aws | | |
| US West (Oregon) | us-west-2 | account-access.us-west-2.api.aws | | |
| Asia Pacific (Mumbai) | ap-south-1 | account-access.ap-south-1.api.aws | | |
| Asia Pacific (Osaka) | ap-northeast-3 | account-access.ap-northeast-3.api.aws | | |
| Asia Pacific (Seoul) | ap-northeast-2 | account-access.ap-northeast-2.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | account-access.ap-southeast-1.api.aws | | |
| Asia Pacific (Sydney) | ap-southeast-2 | account-access.ap-southeast-2.api.aws | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | account-access.ap-northeast-1.api.aws | | |
| Canada (Central) | ca-central-1 | account-access.ca-central-1.api.aws | | |
| Europe (Frankfurt) | eu-central-1 | account-access.eu-central-1.api.aws | | |
| Europe (Ireland) | eu-west-1 | account-access.eu-west-1.api.aws | | |
| Europe (London) | eu-west-2 | account-access.eu-west-2.api.aws | | |
| Europe (Paris) | eu-west-3 | account-access.eu-west-3.api.aws | | |
| Europe (Stockholm) | eu-north-1 | account-access.eu-north-1.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | account-access.sa-east-1.api.aws | | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Group entitlements per role | Each supported Region: 20 | Yes | The maximum number of group entitlements that can be assigned to a single role in an application. |

# Amazon AI Operations endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | aiops.us-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | aiops.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | aiops.us-west-2.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | aiops.eu-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | aiops.eu-central-1.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | aiops.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | aiops.eu-north-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | aiops.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | aiops.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | aiops.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | aiops.ap-northeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | aiops.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | aiops.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | aiops.ap-southeast-7.amazonaws.com | HTTPS |

### Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent active investigations | 2 | No | The maximum number of concurrent investigations with active AI analysis per account per Region. |
| Investigation groups | 1 | No | The maximum number of investigation groups that can be created per account in a Region. |
| Monthly investigations | 150 | No | The maximum number of AI-assist ed investigations per month in this account in the current Region. |

For more information, see CloudWatch Quotas in the *Amazon CloudWatch User Guide*.

## AWS Amplify endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

### Amplify endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | amplify.us-east-2.amazonaws.com | HTTPS | |
| | | amplify.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | amplify.us-east-1.amazonaws.com<br>amplify.us-east-1.api.aws | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | amplify.us-west-1.amazonaws.com<br>amplify.us-west-1.api.aws | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | amplify.us-west-2.amazonaws.com<br>amplify.us-west-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | amplify.ap-east-1.amazonaws.com<br>amplify.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | amplify.ap-south-1.amazonaws.com<br>amplify.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | amplify.ap-northeast-3.amazonaws.com<br>amplify.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | amplify.ap-northeast-2.amazonaws.com<br>amplify.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | amplify.ap-southeast-1.amazonaws.com<br>amplify.ap-southeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | amplify.ap-southeast-2.amazonaws.com | HTTPS |
| | | amplify.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | amplify.ap-northeast-1.amazonaws.com | HTTPS |
| | | amplify.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | amplify.ca-central-1.amazonaws.com | HTTPS |
| | | amplify.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | amplify.eu-central-1.amazonaws.com | HTTPS |
| | | amplify.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | amplify.eu-west-1.amazonaws.com | HTTPS |
| | | amplify.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | amplify.eu-west-2.amazonaws.com | HTTPS |
| | | amplify.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | amplify.eu-south-1.amazonaws.com | HTTPS |
| | | amplify.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | amplify.eu-west-3.amazonaws.com | HTTPS |
| | | amplify.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | amplify.eu-north-1.amazonaws.com | HTTPS |
| | | amplify.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | amplify.me-south-1.amazonaws.com<br>amplify.me-south-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | amplify.sa-east-1.amazonaws.com<br>amplify.sa-east-1.api.aws | HTTPS<br>HTTPS |

## Amplify Studio (backend) endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | amplifybackend.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | amplifybackend.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | amplifybackend.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | amplifybackend.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | amplifybackend.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | amplifybackend.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | amplifybackend.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | amplifybackend.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | amplifybackend.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | amplifybackend.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | amplifybackend.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | amplifybackend.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | amplifybackend.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | amplifybackend.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | amplifybackend.eu-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | amplifybackend.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | amplifybackend.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | amplifybackend.me-south-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | amplifybackend.sa-east-1.amazonaws.com | HTTPS | |

## Amplify Studio (UI Builder) endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | amplifyuibuilder.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | amplifyuibuilder.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | amplifyuibuilder.us-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | amplifyuibuilder.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | amplifyuibuilder.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | amplifyuibuilder.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | amplifyuibuilder.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | amplifyuibuilder.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | amplifyuibuilder.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | amplifyuibuilder.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | amplifyuibuilder.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | amplifyuibuilder.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | amplifyuibuilder.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | amplifyuibuilder.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | amplifyuibuilder.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | amplifyuibuilder.eu-west-3.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | amplifyuibuilder.eu-north-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | amplifyuibuilder.me-south-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | amplifyuibuilder.sa-east-1.amazonaws.com | HTTPS |

## Amplify Service quotas

**Important**

New AWS accounts have reduced apps and concurrent jobs quotas. AWS raises these quotas automatically based on your usage. You can also request a quota increase.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Apps | Each supported Region: 25 | Yes | The maximum number of apps that you can create in AWS Amplify Console in this account in the current Region. |
| Branches per app | Each supported Region: 50 | No | The maximum number of branches per app that you can create in this account in the current Region. |
| Build artifact size | Each supported Region: 5 Gigabytes | No | The maximum size (in GB) of an app build artifact. A build artifact is deployed by AWS Amplify Console after a build. |
| Cache artifact size | Each supported Region: 5 Gigabytes | No | The maximum size (in GB) of a cache artifact. |
| Concurrent jobs | Each supported Region: 5 | Yes | The maximum number of concurrent jobs that you can create in this account in the current Region. |
| Domains per app | Each supported Region: 5 | Yes | The maximum number of domains per app that you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Environment cache artifact size | Each supported Region: 5 Gigabytes | No | The maximum size (in GB) of the environment cache artifact. |
| Manual deploy ZIP file size | Each supported Region: 5 Gigabytes | No | The maximum size (in GB) of a manual deploy ZIP file. |
| Maximum app creations per hour | Each supported Region: 25 | No | The maximum number of apps that you can create in AWS Amplify Console per hour in this account in the current Region. |
| Request tokens per second | Each supported Region: 20,000 | Yes | The maximum number of request tokens per second for an app. Amplify Hosting allocates tokens to requests based on the amount of resources (processing time and data transfer) that they consume. |
| Subdomains per domain | Each supported Region: 50 | No | The maximum number of subdomains per domain that you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Webhooks per app | Each supported Region: 50 | Yes | The maximum number of webhooks per app that you can create in this account in the current Region. |

## Amplify Studio (UI Builder) Service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Component size | Each supported Region: 350 Kilobytes | No | The maximum size (in KB) of the component. |
| Components per app | Each supported Region: 1,000 | No | The maximum number of components that you can create per Amplify app in the current Region. |
| Form size | Each supported Region: 350 Kilobytes | No | The maximum size (in KB) of the form. |
| Forms per app | Each supported Region: 1,000 | No | The maximum number of forms that you can create per Amplify app in the current Region. |
| Theme size | Each supported Region: 350 Kilobytes | No | The maximum size (in KB) of the theme. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Themes per app | Each supported Region: 1,000 | No | The maximum number of themes that you can create per Amplify app in the current Region. |
| View size | Each supported Region: 350 Kilobytes | No | The maximum size (in KB) of the view. |
| Views per app | Each supported Region: 1,000 | No | The maximum number of views that you can create per Amplify app in the current Region. |

# Amazon API Gateway endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Amazon API Gateway includes the API Gateway Control Plane (for creating and managing APIs) and the API Gateway Data Plane (for calling deployed APIs).

The Route 53 Hosted Zone ID column shows the Route 53 Hosted Zone IDs for API Gateway Regional endpoints. Route 53 Hosted Zone IDs are for use with the `execute-api` (API Gateway component service for API execution) domain. For edge-optimized endpoints, the Route 53 Hosted Zone ID is Z2FDTNDATAQYW2 for all Regions.

## Amazon API Gateway control plane

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | apigateway.us-east-2.amazonaws.com<br>apigateway-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | apigateway.us-east-1.amazonaws.com<br>apigateway-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | apigateway.us-west-1.amazonaws.com<br>apigateway-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | apigateway.us-west-2.amazonaws.com<br>apigateway-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | apigateway.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | apigateway.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | apigateway.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | apigateway.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | apigateway.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | apigateway.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | apigateway.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | apigateway.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | apigateway.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | apigateway.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | apigateway.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | apigateway.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | apigateway.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | apigateway.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | apigateway.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | apigateway.ca-central-1.amazonaws.com<br>apigateway-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | apigateway.ca-west-1.amazonaws.com<br>apigateway-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | apigateway.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | apigateway.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | apigateway.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | apigateway.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | apigateway.eu-west-3.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | apigateway.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | apigateway.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | apigateway.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | apigateway.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | apigateway.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | apigateway.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | apigateway.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | apigateway.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | apigateway.us-gov-east-1.amazonaws.com | HTTPS |
| | | apigateway-fips.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | apigateway.us-gov-west-1.amazonaws.com<br>apigateway-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Amazon API Gateway data plane

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | execute-api.us-east-2.amazonaws.com | HTTPS | ZOJJZC49E0EPZ |
| US East (N. Virginia) | us-east-1 | execute-api.us-east-1.amazonaws.com | HTTPS | Z1UJRXOUMOOFQ8 |
| US West (N. California) | us-west-1 | execute-api.us-west-1.amazonaws.com | HTTPS | Z2MUQ32089INYE |
| US West (Oregon) | us-west-2 | execute-api.us-west-2.amazonaws.com | HTTPS | Z2OJLYMUO9EFXC |
| Africa (Cape Town) | af-south-1 | execute-api.af-south-1.amazonaws.com | HTTPS | Z2DHW2332DAMTN |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | execute-api.ap-east-1.amazonaws.com | HTTPS | Z3FD1VL90ND7K5 |
| Asia Pacific (Hyderabad) | ap-south-2 | execute-api.ap-south-2.amazonaws.com | HTTPS | Z0853509Q1135NJ66RUH |
| Asia Pacific (Jakarta) | ap-southeast-3 | execute-api.ap-southeast-3.amazonaws.com | HTTPS | Z10132843TYUYSLUG4HA3 |
| Asia Pacific (Malaysia) | ap-southeast-5 | execute-api.ap-southeast-5.amazonaws.com | HTTPS | Z0314042F0KBUTZ3X5HF |
| Asia Pacific (Melbourne) | ap-southeast-4 | execute-api.ap-southeast-4.amazonaws.com | HTTPS | Z092189423Y7RJK61311D |
| Asia Pacific (Mumbai) | ap-south-1 | execute-api.ap-south-1.amazonaws.com | HTTPS | Z3VO1THU9YC4UR |
| Asia Pacific (Osaka) | ap-northeast-3 | execute-api.ap-northeast-3.amazonaws.com | HTTPS | Z22ILHG95FLSZ2 |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | execute-api.ap-northeast-2.amazonaws.com | HTTPS | Z20JF4UZKIW1U8 |
| Asia Pacific (Singapore) | ap-southeast-1 | execute-api.ap-southeast-1.amazonaws.com | HTTPS | ZL327KTPIQFUL |
| Asia Pacific (Sydney) | ap-southeast-2 | execute-api.ap-southeast-2.amazonaws.com | HTTPS | Z2RPCDW04V8134 |
| Asia Pacific (Taipei) | ap-east-2 | execute-api.ap-east-2.amazonaws.com | HTTPS | Z02909591O7FG9Q56HWB1 |
| Asia Pacific (Thailand) | ap-southeast-7 | execute-api.ap-southeast-7.amazonaws.com | HTTPS | Z048508712PZLK5NKG8R0 |
| Asia Pacific (Tokyo) | ap-northeast-1 | execute-api.ap-northeast-1.amazonaws.com | HTTPS | Z1YSHQZHG15GKL |
| Canada (Central) | ca-central-1 | execute-api.ca-central-1.amazonaws.com | HTTPS | Z19DQILCV0OWEC |
| Canada West (Calgary) | ca-west-1 | execute-api.ca-west-1.amazonaws.com | HTTPS | Z047454934436AWVTG1OQY |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | execute-api.eu-central-1.amazonaws.com | HTTPS | Z1U9ULNL0V5AJ3 |
| Europe (Ireland) | eu-west-1 | execute-api.eu-west-1.amazonaws.com | HTTPS | ZLY8HYME6SFDD |
| Europe (London) | eu-west-2 | execute-api.eu-west-2.amazonaws.com | HTTPS | ZJ5UAJN8Y3Z2Q |
| Europe (Milan) | eu-south-1 | execute-api.eu-south-1.amazonaws.com | HTTPS | Z3BT4WSQ9TDYZV |
| Europe (Paris) | eu-west-3 | execute-api.eu-west-3.amazonaws.com | HTTPS | Z3KY65QIEKYHQQ |
| Europe (Spain) | eu-south-2 | execute-api.eu-south-2.amazonaws.com | HTTPS | Z02499852UI5HEQ5JVWX3 |
| Europe (Stockholm) | eu-north-1 | execute-api.eu-north-1.amazonaws.com | HTTPS | Z3UWIKFBOOGXPP |
| Europe (Zurich) | eu-central-2 | execute-api.eu-central-2.amazonaws.com | HTTPS | Z09222482MK253X48U76H |
| Israel (Tel Aviv) | il-central-1 | execute-api.il-central-1.amazonaws.com | HTTPS | Z07264553HBI44N5X2CKP |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | execute-api.mx-central-1.amazonaws.com | HTTPS | Z00020171WIGL5M88SHRM |
| Middle East (Bahrain) | me-south-1 | execute-api.me-south-1.amazonaws.com | HTTPS | Z20ZBPC0SS8806 |
| Middle East (UAE) | me-central-1 | execute-api.me-central-1.amazonaws.com | HTTPS | Z08780021BKYYY8U0YHTV |
| South America (São Paulo) | sa-east-1 | execute-api.sa-east-1.amazonaws.com | HTTPS | ZCMLWB8V5SYIT |
| AWS GovCloud (US-East) | us-gov-east-1 | execute-api.us-gov-east-1.amazonaws.com | HTTPS | Z3SE9ATJYCRCZJ |
| AWS GovCloud (US-West) | us-gov-west-1 | execute-api.us-gov-west-1.amazonaws.com | HTTPS | Z1K6XKP9SAGWDV |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| API Payload Size | Each supported Region: 10 Megabytes | No | Maximum payload size for non WebSocket API. |
| API Stage throttles in a usage plan | Each supported Region: 20 | Yes | The maximum number of API-stage throttle settings you can create in a usage plan in this account in the current region |
| API keys | Each supported Region: 10,000 | No | The maximum number of API keys that you can create in this account in the current region |
| AWS Lambda authorizer result size | Each supported Region: 8 Kilobytes | No | The maximum size of AWS Lambda authorizer result. |
| Client certificates | Each supported Region: 60 | Yes | The maximum number of certificates that you can associate with this account in the current region |
| Connection duration for WebSocket API | Each supported Region: 7,200 Seconds | No | Maximum duration for WebSocket API connection. |
| Custom Domain Names | Each supported Region: 120 | Yes | The maximum number of Custom domain names that you can create |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | in this account in the current region |
| Domain name access associations | Each supported Region: 100 | Yes | The maximum number of domain name access associations that you can create in this account in the current region |
| Edge API URL Length | Each supported Region: 8,192 | No | Length, in characters, of the URL for an edge-optimized API. |
| Edge-optimized APIs | Each supported Region: 120 | No | The maximum number of edge-optimized APIs that you can create in this account in the current region |
| Maximum API caching TTL | Each supported Region: 3,600 Seconds | No | The maximum API caching TTL you can have in this account in the current region. |
| Maximum Authorizers per Portal | Each supported Region: 1 | No | The maximum number of authorizer per portal |
| Maximum Cached Response Size | Each supported Region: 1,048,576 Bytes | No | Maximum size of cached response in bytes |
| Maximum Combined Header Size | Each supported Region: 10,240 Bytes | No | Maximum combined size of all header values |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum Custom Domain Names per Portal | Each supported Region: 1 | No | The maximum number of custom domain names per portal |
| Maximum Documentation Size per ProductRestEndpointPage | Each supported Region: 32,768 Bytes | No | The maximum documentation size per ProductRestEndpointPage |
| Maximum Iterations In Mapping Template | Each supported Region: 1,000 | No | Maximum number of iterations in a #foreach ... #end loop in mapping templates |
| Maximum Logo Size | Each supported Region: 200 Kilobytes | No | The maximum size of a logo |
| Maximum Size per ProductPage body | Each supported Region: 32,768 Bytes | No | The maximum size per ProductPage body |
| Maximum integration timeout in milliseconds | Each supported Region: 29,000 Milliseconds | Yes | The maximum integration timeout (in milliseconds) allowed for your account in the current region. This limit can only be increased for Regional and private APIs. |
| Maximum resource policy size in bytes | Each supported Region: 8,192 | Yes | The maximum resource policy size in bytes you can have in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Method ARN Length | Each supported Region: 1,600 Bytes | No | ARN length of a method with authorization |
| PortalProducts per Portal | Each supported Region: 200 | No | The maximum number of PortalProducts per Portal |
| PortalProducts per account | Each supported Region: 500 | No | The maximum number of PortalProducts per account |
| Portals per account | Each supported Region: 15 | Yes | The maximum number of Portals per account |
| Private APIs | Each supported Region: 600 | No | The maximum number of private APIs that you can create in this account in the current region |
| Private custom domain names | Each supported Region: 50 | Yes | The maximum number of private custom domain names that you can create in this account in the current region |
| ProductPages per PortalProduct | Each supported Region: 40 | Yes | The maximum number of ProductPages per PortalProduct |
| ProductRestEndpointPages per PortalProduct | Each supported Region: 40 | Yes | The maximum number of ProductRestEndpointPages per PortalProduct |
| Regional API URL Length | Each supported Region: 10,240 | No | Length, in characters, of the URL for a regional API |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Regional APIs | Each supported Region: 600 | No | The maximum number of regional APIs that you can create in this account in the current region |
| Resources/Routes per REST/WebSocket API | Each supported Region: 300 | Yes | The maximum number of resources/routes that you can include in a REST or WebSocket API |
| Routes per HTTP API | Each supported Region: 300 | Yes | The maximum number of routes that you can include in an HTTP API |
| RoutingRules Per Domain Name | Each supported Region: 50 | Yes | The maximum number of RoutingRules per DomainName that you can create in this account in the current region |
| Stage Variable Key Length | Each supported Region: 64 | No | Length, in characters, of the key in a stage variable |
| Stage Variable Value Length | Each supported Region: 512 | No | Length, in characters, of the value in a stage variable |
| Stage variables per stage | Each supported Region: 100 | No | Stage variables per stage |
| Stages per API | Each supported Region: 10 | Yes | The maximum number of stages that you can create for an API |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Subnets per VPC link(V2) | Each supported Region: 10 | Yes | The maximum number of subnets per V2 VPC link in this account in the current Region |
| Tags Per Stage | Each supported Region: 50 | No | Maximum tags per stage. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle burst rate | af-south-1: 1,250<br>ap-east-2: 1,250<br>ap-south-2: 1,250<br>ap-southeast-3: 1,250<br>ap-southeast-4: 1,250<br>ap-southeast-5: 1,250<br>ap-southeast-6: 1,250<br>ap-southeast-7: 1,250<br>ca-west-1: 1,250<br>eu-central-2: 1,250<br>eu-south-1: 1,250<br>eu-south-2: 1,250<br>il-central-1: 1,250<br>mx-central-1: 1,250<br>Each of the other supported Regions: 5,000 | No | The maximum number of additional requests per second (RPS) that you can send in one burst in this account in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate | af-south-1: 2,500<br><br>ap-east-2: 2,500<br><br>ap-south-2: 2,500<br><br>ap-southeast-3: 2,500<br><br>ap-southeast-4: 2,500<br><br>ap-southeast-5: 2,500<br><br>ap-southeast-6: 2,500<br><br>ap-southeast-7: 2,500<br><br>ca-west-1: 2,500<br><br>eu-central-2: 2,500<br><br>eu-south-1: 2,500<br><br>eu-south-2: 2,500<br><br>il-central-1: 2,500<br><br>mx-central-1: 2,500<br><br>Each of the other supported Regions: 10,000 | Yes | The maximum number of requests per second that your APIs can receive in this account in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Usage plans | Each supported Region: 300 | Yes | The maximum number of usage plans that you can create in this account in the current region |
| Usage plans per API key | Each supported Region: 10 | Yes | The maximum number of usage plans that you can associate with an API key |
| VPC links | Each supported Region: 20 | Yes | The maximum number of VPC links that you can create in this account in the current region |
| VPC links(V2) | Each supported Region: 10 | Yes | The maximum number of V2 VPC links that you can create in this account in the current Region |
| WebSocket Idle Connection Timeout | Each supported Region: 600 Seconds | No | WebSocket API idle connection timeout. |
| WebSocket frame size | Each supported Region: 32 Kilobytes | No | Maximum WebSocket frame size. |
| WebSocket message payload size | Each supported Region: 128 Kilobytes | No | Maximum WebSocket message payload size. |
| WebSocket new connections burst rate | Each supported Region: 500 | No | New connections in burst capacity per account (across all WebSocket APIs) per region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| WebSocket new connections rate | Each supported Region: 500 | Yes | New connections per second per account (across all WebSocket APIs) per region |

For more information, see Quotas in Amazon API Gateway in the *API Gateway Developer Guide*.

# AWS AppConfig endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

AWS AppConfig is a capability of AWS Systems Manager. To view endpoints and quotas of other Systems Manager capabilities, see AWS Systems Manager endpoints and quotas.

## Service endpoints

The following sections describe the service endpoints for AWS AppConfig. AWS AppConfig uses *control plane* APIs for setting up and configuring AWS AppConfig applications, environments, configuration profiles, and deployment strategies. AWS AppConfig uses the AWS AppConfig Data service to call *data plane* APIs for retrieving stored configurations.

**Topics**

- Control plane endpoints
- Data plane endpoints

## Control plane endpoints

The following table contains AWS Region-specific endpoints that AWS AppConfig supports for control plane operations. Control plane operations are used for creating, updating, and managing configuration data. For more information, see AWS AppConfig operations in the *AWS AppConfig API Reference*.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | appconfig.us-east-2.amazonaws.com | HTTPS |
| | | appconfig-fips.us-east-2.api.aws | HTTPS |
| | | appconfig-fips.us-east-2.amazonaws.com | HTTPS |
| | | appconfig.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | appconfig.us-east-1.amazonaws.com | HTTPS |
| | | appconfig-fips.us-east-1.api.aws | HTTPS |
| | | appconfig-fips.us-east-1.amazonaws.com | HTTPS |
| | | appconfig.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | appconfig.us-west-1.amazonaws.com | HTTPS |
| | | appconfig-fips.us-west-1.api.aws | HTTPS |
| | | appconfig-fips.us-west-1.amazonaws.com | HTTPS |
| | | appconfig.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | appconfig.us-west-2.amazonaws.com | HTTPS |
| | | appconfig-fips.us-west-2.api.aws | HTTPS |
| | | appconfig-fips.us-west-2.amazonaws.com | HTTPS |
| | | appconfig.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | appconfig.af-south-1.amazonaws.com<br>appconfig.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | appconfig.ap-east-1.amazonaws.com<br>appconfig.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | appconfig.ap-south-2.amazonaws.com<br>appconfig.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | appconfig.ap-southeast-3.amazonaws.com<br>appconfig.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | appconfig.ap-southeast-5.amazonaws.com<br>appconfig.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | appconfig.ap-southeast-4.amazonaws.com<br>appconfig.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | appconfig.ap-south-1.amazonaws.com<br>appconfig.ap-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | appconfig.ap-southeast-6.amazonaws.com<br>appconfig.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | appconfig.ap-northeast-3.amazonaws.com<br>appconfig.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | appconfig.ap-northeast-2.amazonaws.com<br>appconfig.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | appconfig.ap-southeast-1.amazonaws.com<br>appconfig.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | appconfig.ap-southeast-2.amazonaws.com<br>appconfig.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | appconfig.ap-east-2.amazonaws.com<br>appconfig.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | appconfig.ap-southeast-7.amazonaws.com<br>appconfig.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | appconfig.ap-northeast-1.amazonaws.com<br>appconfig.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | appconfig.ca-central-1.amazonaws.com | HTTPS |
| | | appconfig-fips.ca-central-1.api.aws | HTTPS |
| | | appconfig-fips.ca-central-1.amazonaws.com | HTTPS |
| | | appconfig.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | appconfig.ca-west-1.amazonaws.com | HTTPS |
| | | appconfig-fips.ca-west-1.api.aws | HTTPS |
| | | appconfig-fips.ca-west-1.amazonaws.com | HTTPS |
| | | appconfig.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | appconfig.eu-central-1.amazonaws.com | HTTPS |
| | | appconfig.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | appconfig.eu-west-1.amazonaws.com | HTTPS |
| | | appconfig.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | appconfig.eu-west-2.amazonaws.com | HTTPS |
| | | appconfig.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | appconfig.eu-south-1.amazonaws.com | HTTPS |
| | | appconfig.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | appconfig.eu-west-3.amazonaws.com | HTTPS |
| | | appconfig.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | appconfig.eu-south-2.amazonaws.com | HTTPS |
| | | appconfig.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | appconfig.eu-north-1.amazonaws.com | HTTPS |
| | | appconfig.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | appconfig.eu-central-2.amazonaws.com | HTTPS |
| | | appconfig.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | appconfig.il-central-1.amazonaws.com | HTTPS |
| | | appconfig.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | appconfig.mx-central-1.amazonaws.com | HTTPS |
| | | appconfig.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | appconfig.me-south-1.amazonaws.com | HTTPS |
| | | appconfig.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | appconfig.me-central-1.amazonaws.com | HTTPS |
| | | appconfig.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | appconfig.sa-east-1.amazonaws.com | HTTPS |
| | | appconfig.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | appconfig.us-gov-east-1.amazonaws.com | HTTPS |
| | | appconfig-fips.us-gov-east-1.api.aws | HTTPS |
| | | appconfig-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | appconfig.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | appconfig.us-gov-west-1.amazonaws.com | HTTPS |
| | | appconfig-fips.us-gov-west-1.api.aws | HTTPS |
| | | appconfig-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | appconfig.us-gov-west-1.api.aws | HTTPS |

## Data plane endpoints

The following table contains AWS Region-specific endpoints that AWS AppConfig Data supports for data plane operations. Data plane operations are used for retrieving configuration data. For more information, see AWS AppConfig Data operations in the *AWS AppConfig API Reference*.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | appconfigdata.us-east-2.amazonaws.com | HTTPS |
| | | appconfigdata-fips.us-east-2.api.aws | HTTPS |
| | | appconfigdata-fips.us-east-2.amazonaws.com | HTTPS |
| | | appconfigdata.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | appconfigdata.us-east-1.amazonaws.com | HTTPS |
| | | appconfigdata-fips.us-east-1.api.aws | HTTPS |
| | | appconfigdata-fips.us-east-1.amazonaws.com | HTTPS |
| | | appconfigdata.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | appconfigdata.us-west-1.amazonaws.com | HTTPS |
| | | appconfigdata-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | appconfigdata-fips.us-west-1.amazonaws.com | HTTPS |
| | | appconfigdata.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | appconfigdata.us-west-2.amazonaws.com | HTTPS |
| | | appconfigdata-fips.us-west-2.api.aws | HTTPS |
| | | appconfigdata-fips.us-west-2.amazonaws.com | HTTPS |
| | | appconfigdata.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | appconfigdata.af-south-1.amazonaws.com | HTTPS |
| | | appconfigdata.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | appconfigdata.ap-east-1.amazonaws.com | HTTPS |
| | | appconfigdata.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | appconfigdata.ap-south-2.amazonaws.com | HTTPS |
| | | appconfigdata.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | appconfigdata.ap-southeast-3.amazonaws.com | HTTPS |
| | | appconfigdata.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | appconfigdata.ap-southeast-5.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | appconfigdata.ap-southeast-4.amazonaws.com<br><br>appconfigdata.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | appconfigdata.ap-south-1.amazonaws.com<br><br>appconfigdata.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | appconfigdata.ap-southeast-6.amazonaws.com<br><br>appconfigdata.ap-southeast-6.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | appconfigdata.ap-northeast-3.amazonaws.com<br><br>appconfigdata.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | appconfigdata.ap-northeast-2.amazonaws.com<br><br>appconfigdata.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | appconfigdata.ap-southeast-1.amazonaws.com<br><br>appconfigdata.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | appconfigdata.ap-southeast-2.amazonaws.com<br><br>appconfigdata.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | appconfigdata.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | appconfigdata.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | appconfigdata.ap-northeast-1.amazonaws.com<br>appconfigdata.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | appconfigdata.ca-central-1.amazonaws.com<br>appconfigdata-fips.ca-central-1.api.aws<br>appconfigdata-fips.ca-central-1.amazonaws.com<br>appconfigdata.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | appconfigdata.ca-west-1.amazonaws.com<br>appconfigdata-fips.ca-west-1.api.aws<br>appconfigdata-fips.ca-west-1.amazonaws.com<br>appconfigdata.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | appconfigdata.eu-central-1.amazonaws.com<br>appconfigdata.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | appconfigdata.eu-west-1.amazonaws.com<br>appconfigdata.eu-west-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | appconfigdata.eu-west-2.amazonaws.com<br>appconfigdata.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | appconfigdata.eu-south-1.amazonaws.com<br>appconfigdata.eu-south-1.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | appconfigdata.eu-west-3.amazonaws.com<br>appconfigdata.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | appconfigdata.eu-south-2.amazonaws.com<br>appconfigdata.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | appconfigdata.eu-north-1.amazonaws.com<br>appconfigdata.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | appconfigdata.eu-central-2.amazonaws.com<br>appconfigdata.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | appconfigdata.il-central-1.amazonaws.com<br>appconfigdata.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | appconfigdata.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | appconfigdata.me-south-1.amazonaws.com<br>appconfigdata.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | appconfigdata.me-central-1.amazonaws.com<br>appconfigdata.me-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | appconfigdata.sa-east-1.amazonaws.com<br>appconfigdata.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | appconfigdata.us-gov-east-1.amazonaws.com<br>appconfigdata.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | appconfigdata.us-gov-west-1.amazonaws.com<br>appconfigdata.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Configuration size limit in AWS AppConfig hosted configuration store | Each supported Region: 2,048 Kilobytes | Yes | The current maximum approval limit of a hosted configuration is 4,000 Kilobytes per version. There is no additional cost to use hosted configurations. To reduce configuration size, remove unnecessary configuration data, or segment data across multiple configuration profiles. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Deployment size limit | Each supported Region: 2,048 Kilobytes | Yes | The current maximum approval limit of an AWS AppConfig deployment is 4,000 Kilobytes per deployment. To reduce configuration size, remove unnecessary configuration data, or segment configuration data across multiple configuration profiles. |
| Maximum number of applications | Each supported Region: 100 | Yes | An application in AWS AppConfig is a logical unit of code (a namespace) that provides capabilities for your customers. Limit increase requests will be automatically approved for values less than or equal to 300. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of configuration profiles per application | Each supported Region: 100 | Yes | A configuration profile contains metadata about a particular set of configuration data used by your application. To reduce the number of configuration profiles required per application, split configuration profiles across multiple applications. Limit increase requests will be automatically approved for values less than or equal to 500. |
| Maximum number of deployment strategies | Each supported Region: 20 | Yes | A deployment strategy defines how configuration deploys, or rolls out, across a collection of instances within a specific application environment. Limit increase requests will be automatically approved for values less than or equal to 100. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of environments per application | Each supported Region: 20 | Yes | An environment corresponds to a grouping of instances associated with an application. Environments examples include stages, such as beta and prod, or application subcomponents, such as web, mobile, and service. Limit increase requests will be automatically approved for values less than or equal to 100. |
| Maximum number of experiment definitions per application | Each supported Region: 50 | Yes | An experiment definition configures A/B testing for a feature flag within an application. To reduce the number of experiment definitions required per application, archive unused experiment definitions. Limit increase requests will be automatically approved for values less than or equal to 500. |

## Control plane default limits

| API name | Transactions per second | Adjustable |
|---|---|---|
| Create* | 10 | No |
| Delete* | 10 | No |
| Get* | 100 | No |
| List* | 10 | No |
| StartDeployment | 10 | No |
| StopDeployment | 10 | No |
| TagResource | 20 | No |
| UntagResource | 20 | No |
| Update* | 10 | No |
| ValidateConfiguration | 10 | No |

## Data plane default limits

| Action | API limit | Adjustable |
|---|---|---|
| GetConfiguration (deprecated) | 500 TPS | No |
| GetLatestConfiguration | 1,000 TPS | Yes |
| StartConfigurationSession | 500 TPS | Yes |
| Configurations received | 1 million (burst) per day if not using the AWS AppConfig agent | Yes |

To request an increase for `GetLatestConfiguration`, `StartConfigurationSession`, or `Configurations received` contact Support. To improve performance, availability, and reduce costs, we recommended you cache configurations locally when using AWS AppConfig. AWS AppConfig Agent caches configurations on your behalf.

# AWS AppFabric endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | appfabric.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | appfabric.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | appfabric.eu-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Application authorizations | Each supported Region: 50 | No | The maximum number of application authorizations that you can create in an account in the current AWS Region. |
| Application bundles | Each supported Region: 1 | No | The maximum number of application bundles that you can create in an account in the current AWS Region. |
| Ingestion destinations | Each supported Region: 5 | No | The maximum number of ingestion destinations that you can create per ingestion in an account in the current AWS Region. |
| Ingestions | Each supported Region: 50 | No | The maximum number of ingestions that you can create in an account in the current AWS Region. |

# Amazon Application Recovery Controller (ARC) endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Zonal shift

Zonal shift in ARC is available in all AWS Regions, including the Beijing and Ningxia Regions and AWS GovCloud (US). The Region switch, routing control, and readiness check capabilities of the ARC service are not available in the Beijing and Ningxia Regions nor in AWS GovCloud (US).

For the ARC **Zonal Shift API**, including API operations for zonal autoshift in the Regions that include the zonal autoshift capability, use the following endpoints.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | arc-zonal-shift.us-east-2.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.us-east-2.api.aws | HTTPS |
| | | arc-zonal-shift.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | arc-zonal-shift.us-east-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.us-east-1.api.aws | HTTPS |
| | | arc-zonal-shift.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | arc-zonal-shift.us-west-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.us-west-1.api.aws | HTTPS |
| | | arc-zonal-shift.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | arc-zonal-shift.us-west-2.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.us-west-2.api.aws | HTTPS |
| | | arc-zonal-shift.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | arc-zonal-shift.af-south-1.amazonaws.com<br>arc-zonal-shift.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | arc-zonal-shift.ap-east-1.amazonaws.com<br>arc-zonal-shift.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | arc-zonal-shift.ap-south-2.amazonaws.com<br>arc-zonal-shift.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | arc-zonal-shift.ap-southeast-3.amazonaws.com<br>arc-zonal-shift.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | arc-zonal-shift.ap-southeast-5.amazonaws.com<br>arc-zonal-shift.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | arc-zonal-shift.ap-southeast-4.amazonaws.com<br>arc-zonal-shift.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | arc-zonal-shift.ap-south-1.amazonaws.com<br>arc-zonal-shift.ap-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | arc-zonal-shift.ap-southeast-6.amazonaws.com<br>arc-zonal-shift.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | arc-zonal-shift.ap-northeast-3.amazonaws.com<br>arc-zonal-shift.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | arc-zonal-shift.ap-northeast-2.amazonaws.com<br>arc-zonal-shift.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | arc-zonal-shift.ap-southeast-1.amazonaws.com<br>arc-zonal-shift.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | arc-zonal-shift.ap-southeast-2.amazonaws.com<br>arc-zonal-shift.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | arc-zonal-shift.ap-east-2.amazonaws.com<br>arc-zonal-shift.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | arc-zonal-shift.ap-southeast-7.amazonaws.com<br>arc-zonal-shift.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | arc-zonal-shift.ap-northeast-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | arc-zonal-shift.ca-central-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.ca-central-1.api.aws | HTTPS |
| | | arc-zonal-shift.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | arc-zonal-shift.ca-west-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.ca-west-1.api.aws | HTTPS |
| | | arc-zonal-shift.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | arc-zonal-shift.eu-central-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | arc-zonal-shift.eu-west-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | arc-zonal-shift.eu-west-2.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | arc-zonal-shift.eu-south-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | arc-zonal-shift.eu-west-3.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | arc-zonal-shift.eu-south-2.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | arc-zonal-shift.eu-north-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | arc-zonal-shift.eu-central-2.amazonaws.com | HTTPS |
| | | arc-zonal-shift.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | arc-zonal-shift.il-central-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | arc-zonal-shift.mx-central-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | arc-zonal-shift.me-south-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | arc-zonal-shift.me-central-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | arc-zonal-shift.sa-east-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | arc-zonal-shift.us-gov-east-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.us-gov-east-1.api.aws | HTTPS |
| | | arc-zonal-shift.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | arc-zonal-shift.us-gov-west-1.amazonaws.com | HTTPS |
| | | arc-zonal-shift-fips.us-gov-west-1.api.aws | HTTPS |
| | | arc-zonal-shift.us-gov-west-1.api.aws | HTTPS |

## Regional features

For the ARC **Region switch API**, use the following endpoints.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | arc-region-switch.us-east-2.api.aws | HTTPS |
| | | arc-region-switch-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | arc-region-switch.us-east-1.api.aws | HTTPS |
| | | arc-region-switch-control-plane-fips.us-east-1.api.aws | HTTPS |
| | | arc-region-switch-fips.us-east-1.api.aws | HTTPS |
| | | arc-region-switch-control-plane.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | arc-region-switch.us-west-1.api.aws | HTTPS |
| | | arc-region-switch-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | arc-region-switch.us-west-2.api.aws | HTTPS |
| | | arc-region-switch-fips.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | arc-region-switch.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | arc-region-switch.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | arc-region-switch.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | arc-region-switch.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | arc-region-switch.ap-southeast-5.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | arc-region-switch.ap-southeast-4.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | arc-region-switch.ap-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | arc-region-switch.ap-southeast-6.api.aws | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | arc-region-switch.ap-northeast-3.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | arc-region-switch.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | arc-region-switch.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | arc-region-switch.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | arc-region-switch.ap-east-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | arc-region-switch.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | arc-region-switch.ap-northeast-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | arc-region-switch.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | arc-region-switch.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | arc-region-switch.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | arc-region-switch.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | arc-region-switch.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | arc-region-switch.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | arc-region-switch.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | arc-region-switch.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | arc-region-switch.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | arc-region-switch.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | arc-region-switch.il-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | arc-region-switch.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | arc-region-switch.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | arc-region-switch.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | arc-region-switch.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | arc-region-switch.us-gov-east-1.api.aws<br>arc-region-switch-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | arc-region-switch.us-gov-west-1.api.aws<br>arc-region-switch-control-plane-fips.us-gov-west-1.api.aws<br>arc-region-switch-fips.us-gov-west-1.api.aws<br>arc-region-switch-control-plane.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

When you use the AWS CLI or SDKs to submit requests with ARC **Recovery Readiness API** (for readiness checks), **Recovery Control Configuration API** or **Recovery Cluster API** (for routing control), you must specify the AWS Region as `us-west-2`.

For the ARC **Recovery Readiness API** (for readiness checks) or **Recovery Control Configuration API**, use the following endpoints, respectively.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) Region | us-west-2 | route53-recovery-readiness.amazonaws.com | HTTPS |
| US West (Oregon) Region | us-west-2 | route53-recovery-control-config.amazonaws.com | HTTPS |

For the ARC **Recovery Cluster API**, in addition to specifying the Region as `us-west-2`, you also must specify one of your five Regional cluster endpoints. The endpoint that you specify must target the ARC cluster that hosts the routing controls that you want to get or update the state for.

ARC creates endpoints for each cluster in the following five Regions: US East (N. Virginia) (us-east-1), Europe (Ireland) (eu-west-1), US West (Oregon) (us-west-2), Asia Pacific (Tokyo) (ap-northeast-1), and Asia Pacific (Sydney) (ap-southeast-2). Routing Controls provide five regional endpoints to ensure high availability, even in the face of failures. To achieve their full resilience, it's important to have retry logic that can use all five endpoints as necessary. To learn more, see Get and update routing control states using the API and Best practices for Amazon Application Recovery Controller (ARC) in the Amazon Route 53 Application Recovery Controller Developer Guide.

The following are examples of the Regional cluster endpoints for routing control in ARC.

| Endpoint | Region |
|---|---|
| https://aaaaaaaa.route53-recovery-cluster.eu-west-1.amazonaws.com | eu-west-1 |
| https://bbbbbbb.route53-recovery-cluster.ap-northeast-1.amazonaws.com | ap-northeast-1 |
| https://ccccccc.route53-recovery-cluster.us-west-2.amazonaws.com | us-west-2 |
| https://ddddddd.route53-recovery-cluster.us-east-1.amazonaws.com | us-east-1 |
| https://eeeeeee.route53-recovery-cluster.ap-southeast-2.amazonaws.com | ap-southeast-2 |

## Service quotas

For more information, see Quotas in Amazon Application Recovery Controller (ARC) in the *Amazon Application Recovery Controller (ARC) Developer Guide*.

# AWS App Mesh endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | appmesh.us-east-2.amazonaws.com | HTTPS |
| | | appmesh-envoy-management-fips.us-east-2.amazonaws.com | HTTPS |
| | | appmesh.us-east-2.api.aws | HTTPS |
| | | appmesh-envoy-management-fips.us-east-2.api.aws | HTTPS |
| | | appmesh-envoy-management.us-east-2.amazonaws.com | HTTPS |
| | | appmesh-envoy-management.us-east-2.api.aws | HTTPS |
| | | appmesh-fips.us-east-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | appmesh-fips.us-east-2.amazonaws.com | | |
| US East (N. Virginia) | us-east-1 | appmesh.us-east-1.amazonaws.com<br>appmesh-envoy-management.us-east-1.api.aws<br>appmesh-fips.us-east-1.amazonaws.com<br>appmesh-envoy-management-fips.us-east-1.api.aws<br>appmesh-envoy-management.us-east-1.amazonaws.com<br>appmesh.us-east-1.api.aws<br>appmesh-fips.us-east-1.api.aws<br>appmesh-envoy-management-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (N. California) | us-west-1 | appmesh.us-west-1.amazonaws.com | HTTPS | |
| | | appmesh-envoy-management.us-west-1.amazonaws.com | HTTPS | |
| | | appmesh.us-west-1.api.aws | HTTPS | |
| | | appmesh-fips.us-west-1.api.aws | HTTPS | |
| | | appmesh-fips.us-west-1.amazonaws.com | HTTPS | |
| | | appmesh-envoy-management.us-west-1.api.aws | HTTPS | |
| | | appmesh-envoy-management-fips.us-west-1.amazonaws.com | HTTPS | |
| | | appmesh-envoy-management-fips.us-west-1.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | appmesh.us-west-2.amazonaws.com<br>appmesh-envoy-management.us-west-2.amazonaws.com<br>appmesh-envoy-management-fips.us-west-2.amazonaws.com<br>appmesh-fips.us-west-2.amazonaws.com<br>appmesh.us-west-2.api.aws<br>appmesh-fips.us-west-2.api.aws<br>appmesh-envoy-management-fips.us-west-2.api.aws<br>appmesh-envoy-management.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | appmesh.af-south-1.amazonaws.com<br>appmesh-envoy-management.af-south-1.amazonaws.com<br>appmesh.af-south-1.api.aws<br>appmesh-envoy-management.af-south-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | appmesh.ap-east-1.amazonaws.com<br>appmesh.ap-east-1.api.aws<br>appmesh-envoy-management.ap-east-1.api.aws<br>appmesh-envoy-management.ap-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | appmesh.ap-southeast-3.amazonaws.com<br>appmesh-envoy-management.ap-southeast-3.amazonaws.com<br>appmesh.ap-southeast-3.api.aws<br>appmesh-envoy-management.ap-southeast-3.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | appmesh.ap-south-1.amazonaws.com<br>appmesh.ap-south-1.api.aws<br>appmesh-envoy-management.ap-south-1.amazonaws.com<br>appmesh-envoy-management.ap-south-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | appmesh.ap-northeast-3.amazonaws.com<br>appmesh-envoy-management.ap-northeast-3.amazonaws.com<br>appmesh-envoy-management.ap-northeast-3.api.aws<br>appmesh.ap-northeast-3.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | appmesh.ap-northeast-2.amazonaws.com<br>appmesh-envoy-management.ap-northeast-2.api.aws<br>appmesh-envoy-management.ap-northeast-2.amazonaws.com<br>appmesh.ap-northeast-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | appmesh.ap-southeast-1.amazonaws.com<br>appmesh-envoy-management.ap-southeast-1.api.aws<br>appmesh.ap-southeast-1.api.aws<br>appmesh-envoy-management.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | appmesh.ap-southeast-2.amazonaws.com<br>appmesh-envoy-management.ap-southeast-2.api.aws<br>appmesh-envoy-management.ap-southeast-2.amazonaws.com<br>appmesh.ap-southeast-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | appmesh.ap-northeast-1.amazonaws.com<br>appmesh-envoy-management.ap-northeast-1.api.aws<br>appmesh-envoy-management.ap-northeast-1.amazonaws.com<br>appmesh.ap-northeast-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | appmesh.ca-central-1.amazonaws.com<br>appmesh-envoy-management-fips.ca-central-1.api.aws<br>appmesh.ca-central-1.api.aws<br>appmesh-envoy-management.ca-central-1.api.aws<br>appmesh-fips.ca-central-1.amazonaws.com<br>appmesh-envoy-management.ca-central-1.amazonaws.com<br>appmesh-fips.ca-central-1.api.aws<br>appmesh-envoy-management-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | appmesh.eu-central-1.amazonaws.com<br>appmesh.eu-central-1.api.aws<br>appmesh-envoy-management.eu-central-1.amazonaws.com<br>appmesh-envoy-management.eu-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | appmesh.eu-west-1.amazonaws.com<br>appmesh-envoy-management.eu-west-1.api.aws<br>appmesh-envoy-management.eu-west-1.amazonaws.com<br>appmesh.eu-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | appmesh.eu-west-2.amazonaws.com<br>appmesh.eu-west-2.api.aws<br>appmesh-envoy-management.eu-west-2.amazonaws.com<br>appmesh-envoy-management.eu-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | appmesh.eu-south-1.amazonaws.com<br>appmesh-envoy-management.eu-south-1.api.aws<br>appmesh-envoy-management.eu-south-1.amazonaws.com<br>appmesh.eu-south-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | appmesh.eu-west-3.amazonaws.com<br>appmesh-envoy-management.eu-west-3.amazonaws.com<br>appmesh-envoy-management.eu-west-3.api.aws<br>appmesh.eu-west-3.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Spain) | eu-south-2 | appmesh.eu-south-2.amazonaws.com<br>appmesh-envoy-management.eu-south-2.api.aws<br>appmesh.eu-south-2.api.aws<br>appmesh-envoy-management.eu-south-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | appmesh.eu-north-1.amazonaws.com<br>appmesh.eu-north-1.api.aws<br>appmesh-envoy-management.eu-north-1.api.aws<br>appmesh-envoy-management.eu-north-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | appmesh.eu-central-2.amazonaws.com<br>appmesh-envoy-management.eu-central-2.amazonaws.com<br>appmesh-envoy-management.eu-central-2.api.aws<br>appmesh.eu-central-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | appmesh.il-central-1.amazonaws.com<br>appmesh-envoy-management.il-central-1.amazonaws.com<br>appmesh-envoy-management.il-central-1.api.aws<br>appmesh.il-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | appmesh.me-south-1.amazonaws.com<br>appmesh-envoy-management.me-south-1.api.aws<br>appmesh-envoy-management.me-south-1.amazonaws.com<br>appmesh.me-south-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | appmesh.sa-east-1.amazonaws.com<br>appmesh.sa-east-1.api.aws<br>appmesh-envoy-management.sa-east-1.api.aws<br>appmesh-envoy-management.sa-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Backends per virtual node | Each supported Region: 50 | No | Number of backends per virtual node |
| Connected Envoy processes per virtual gateway | Each supported Region: 50 | Yes | Number of concurrently connected Envoy processes per virtual gateway |
| Connected Envoy processes per virtual node | Each supported Region: 50 | Yes | Number of concurrently connected Envoy processes per virtual node |
| Gateway routes per virtual gateway | Each supported Region: 10 | Yes | Number of gateway routes per virtual gateway |
| Meshes per account | Each supported Region: 15 | Yes | Number of meshes per account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Routes per virtual router | Each supported Region: 50 | Yes | Number of routes per virtual router |
| Virtual gateways per mesh | Each supported Region: 3 | Yes | Number of virtual gateways per mesh |
| Virtual nodes per mesh | Each supported Region: 200 | Yes | Number of virtual nodes per mesh |
| Virtual routers per mesh | Each supported Region: 200 | Yes | Number of virtual routers per mesh |
| Virtual services per mesh | Each supported Region: 200 | Yes | Number of virtual services per mesh |
| Weighted targets per route | Each supported Region: 10 | No | Number of weighted targets per route |

# AWS App Runner endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | apprunner.us-east-2.amazonaws.com<br>apprunner-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | Z0224347AD7KVHMLOX31 |
| US East (N. Virginia) | us-east-1 | apprunner.us-east-1.amazonaws.com<br>apprunner-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | Z01915732ZBZKC8D32TPT |
| US West (Oregon) | us-west-2 | apprunner.us-west-2.amazonaws.com<br>apprunner-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | Z02243383FTQ64HJ5772Q |
| Asia Pacific (Mumbai) | ap-south-1 | apprunner.ap-south-1.amazonaws.com | HTTPS | Z00855883LBHKTIC4ODF2 |
| Asia Pacific (Singapore) | ap-southeast-1 | apprunner.ap-southeast-1.amazonaws.com | HTTPS | Z09819469CZ3KQ8PWMCL |
| Asia Pacific (Sydney) | ap-southeast-2 | apprunner.ap-southeast-2.amazonaws.com | HTTPS | Z03657752RA8799S0TI5I |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | apprunner.ap-northeast-1.amazonaws.com | HTTPS | Z08491812XW6IPYLR6CCA |
| Europe (Frankfurt) | eu-central-1 | apprunner.eu-central-1.amazonaws.com | HTTPS | Z0334911C2FDI2Q9M4FZ |
| Europe (Ireland) | eu-west-1 | apprunner.eu-west-1.amazonaws.com | HTTPS | Z087551914Z2PCAU0QHMW |
| Europe (London) | eu-west-2 | apprunner.eu-west-2.amazonaws.com | HTTPS | Z098228427VC6B3IX76ON |
| Europe (Paris) | eu-west-3 | apprunner.eu-west-3.amazonaws.com | HTTPS | Z087117439MBKHYM69QS6 |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Auto scaling configurations | Each supported Region: 10 | Yes | The maximum number of auto scaling configurations that you can create in this account in the current Region. You can use a single auto scaling |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | configuration in multiple services. |
| Connections | Each supported Region: 10 | Yes | The maximum number of connections that you can create in this account in the current Region. You can use a single connection in multiple services. |
| Observability configurations | Each supported Region: 10 | Yes | The maximum number of observability configurations that you can create in this account in the current Region. You can use a single observability configuration in multiple services. |
| Services | Each supported Region: 30 | Yes | The maximum number of services that you can create in this account in the current Region. |
| VPC connectors | Each supported Region: 10 | Yes | The maximum number of VPC connectors that you can create in this account in the current Region. You can use a single VPC connector in multiple services. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| VPC ingress connections | Each supported Region: 1 | No | The maximum number of VPC Ingress Connections that you can create for a single Service. |

# Amazon AppFlow endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

You can't use IP allow listing in your Amazon S3 bucket policy to deny access to any other IP addresses besides Amazon AppFlow IP addresses. This is because Amazon AppFlow uses a VPC endpoint when placing data in your Amazon S3 buckets.

For more information about the IP addresses used by Amazon AppFlow, see AWS IP address ranges in the *Amazon Web Services General Reference*.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | appflow.us-east-2.amazonaws.com | HTTPS | |
| | | appflow-fips.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | appflow.us-east-1.amazonaws.com<br>appflow-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | appflow.us-west-1.amazonaws.com<br>appflow-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | appflow.us-west-2.amazonaws.com<br>appflow-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | appflow.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | appflow.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | appflow.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | appflow.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | appflow.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | appflow.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | appflow.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | appflow.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | appflow.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | appflow.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | appflow.eu-west-3.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | appflow.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amazon AppFlow flow run size | Each supported Region: 100 Gigabytes | No | The maximum size (in GB) of data that Amazon AppFlow can process per |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | flow run in this account in the current Region. |
| Amazon EventBridge event size | Each supported Region: 256 Kilobytes | No | The maximum size (in KB) of events that Amazon EventBridge can process per flow run in this account in the current Region. If your event exceeds this size, Amazon AppFlow publishes a summary event with a pointer to the Amazon S3 bucket where you can get the full event. |
| Amplitude flow run size | Each supported Region: 25 Megabytes | No | The maximum size (in MB) of data that can be processed per flow run when when using Amplitude as a source, in this account in the current Region. |
| Concurrent flow runs | Each supported Region: 1,000 | Yes | The maximum number of concurrent flow runs that you can perform at any time, in this account in the current Region. |
| Connector profiles | Each supported Region: 100 | Yes | The maximum number of connector profiles in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Google Analytics dimensions | Each supported Region: 9 | No | The maximum number of Google Analytics dimensions that can be processed per flow run in this account in the current Region. |
| Google Analytics metrics | Each supported Region: 10 | No | The maximum number of Google Analytics metrics that can be processed per flow run in this account in the current Region. |
| Marketo flow run size | Each supported Region: 20 Megabytes | No | The maximum size (in MB) of data that can be processed per flow run when using Marketo as a source, in this account in the current Region. |
| Monthly flow runs | Each supported Region: 10,000,000 | Yes | The maximum number of flow runs you can perform per month in this account in the current Region. |
| Rate of Amazon AppFlow flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that Amazon AppFlow can run per minute in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of Amazon S3 flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Amazon S3 as a source, in this account in the current Region. |
| Rate of Amplitude flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per day when using Amplitude as a source, in this account in the current Region. |
| Rate of Datadog flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Datadog as a source, in this account in the current Region. |
| Rate of Dynatrace flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Dynatrace as a source, in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of Google Analytics flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per day when using Google Analytics as a source, in this account in the current Region. |
| Rate of Infor Nexus flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Infor Nexus as a source, in this account in the current Region. |
| Rate of Marketo flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per hour when using Marketo as a source, in this account in the current Region. |
| Rate of Salesforce Pardot flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Salesforce Pardot as a source, in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of Salesforce flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Salesforce as a source, in this account in the current Region. |
| Rate of ServiceNow flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using ServiceNow as a source, in this account in the current Region. |
| Rate of Singular flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per hour when using Singular as a source, in this account in the current Region. |
| Rate of Slack flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Slack as a source, in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of TrendMicro flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per hour when using TrendMicro as a source, in this account in the current Region. |
| Rate of Veeva flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Veeva as a source, in this account in the current Region. |
| Rate of Zendesk flow runs | Each supported Region: 1 | No | The maximum number of schedule-triggered flows that you can run per minute when using Zendesk as a source, in this account in the current Region. |
| Salesforce event size | Each supported Region: 1 Megabytes | No | The maximum size (in MB) of events from Salesforce that can be processed per flow run in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Salesforce flow run data export size | Each supported Region: 500 Megabytes | No | The maximum size (in MB) of records that you can insert, update, or upsert into Salesforce per flow run. If your source is Amazon S3, each CSV file cannot exceed 125 MB in size. However, you can drop multiple CSV files into the source bucket or folder, and Amazon AppFlow will transfer all the data to Salesforce in a single flow run. |
| Salesforce flow run data import size | Each supported Region: 15 Gigabytes | No | The maximum size (in GB) of data that Salesforce can import per flow run in this account in the current Region. |
| ServiceNow records | Each supported Region: 100,000 | No | The maximum number of ServiceNow records that can be processed per flow run in this account in the current Region. |
| Total flows | Each supported Region: 1,000 | Yes | The maximum number of flows that you can have in this account in the current Region. |

For more information, see Quotas for Amazon AppFlow in the *Amazon AppFlow User Guide*.

# Application Auto Scaling endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | application-autoscaling.us-east-2.amazonaws.com | HTTP and HTTPS | |
| US East (N. Virginia) | us-east-1 | application-autoscaling.us-east-1.amazonaws.com | HTTP and HTTPS | |
| US West (N. California) | us-west-1 | application-autoscaling.us-west-1.amazonaws.com | HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | application-autoscaling.us-west-2.amazonaws.com | HTTP and HTTPS | |
| Africa (Cape Town) | af-south-1 | application-autoscaling.af-south-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | application-autoscaling.ap-east-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | application-autoscaling.ap-south-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | application-autoscaling.ap-southeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | application-autoscaling.ap-southeast-5.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | application-autoscaling.ap-southeast-4.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | application-autoscaling.ap-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | application-autoscaling.ap-southeast-6.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | application-autoscaling.ap-northeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | application-autoscaling.ap-northeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | application-autoscaling.ap-southeast-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | application-autoscaling.ap-southeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | application-autoscaling.ap-east-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | application-autoscaling.ap-southeast-7.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | application-autoscaling.ap-northeast-1.amazonaws.com | HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | application-autoscaling.ca-central-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | application-autoscaling.ca-west-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | application-autoscaling.eu-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | application-autoscaling.eu-west-1.amazonaws.com | HTTP and HTTPS | |
| Europe (London) | eu-west-2 | application-autoscaling.eu-west-2.amazonaws.com | HTTP and HTTPS | |
| Europe (Milan) | eu-south-1 | application-autoscaling.eu-south-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | application-autoscaling.eu-west-3.amazonaws.com | HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | application-autoscaling.eu-south-2.amazonaws.com | HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | application-autoscaling.eu-north-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Zurich) | eu-central-2 | application-autoscaling.eu-central-2.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | application-autoscaling.il-central-1.amazonaws.com | HTTP and HTTPS |
| Mexico (Central) | mx-central-1 | application-autoscaling.mx-central-1.amazonaws.com | HTTP and HTTPS |
| Middle East (Bahrain) | me-south-1 | application-autoscaling.me-south-1.amazonaws.com | HTTP and HTTPS |
| Middle East (UAE) | me-central-1 | application-autoscaling.me-central-1.amazonaws.com | HTTP and HTTPS |
| South America (São Paulo) | sa-east-1 | application-autoscaling.sa-east-1.amazonaws.com | HTTP and HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | application-autoscaling.us-gov-east-1.amazonaws.com | HTTP and HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | application-autoscaling.us-gov-west-1.amazonaws.com | HTTP and HTTPS |

# Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Scalable targets for Amazon ElastiCache | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the ElastiCache namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for Amazon Keyspaces | Each supported Region: 1,500 | Yes | The maximum number of scalable targets that you can register for the Cassandra namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for Amazon MSK | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the Kafka namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Scalable targets for AppStream | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the AppStream namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for Comprehend | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the Comprehend namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for DynamoDB | Each supported Region: 5,000 | Yes | The maximum number of scalable targets that you can register for the DynamoDB namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Scalable targets for EC2 | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the EC2 namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for ECS | Each supported Region: 3,000 | Yes | The maximum number of scalable targets that you can register for the ECS namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for EMR | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the Elastic MapReduce (EMR) namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Scalable targets for Lambda | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the Lambda namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for RDS | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the RDS namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scalable targets for SageMaker | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the SageMaker namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Scalable targets for custom resources | Each supported Region: 500 | Yes | The maximum number of scalable targets that you can register for the custom resource namespace in this account in the current Region. A scalable target identifies the resource that Application Auto Scaling can scale. |
| Scaling policies per scalable target | Each supported Region: 50 | No | The maximum number of scaling policies per scalable target. |
| Scheduled actions per scalable target | Each supported Region: 200 | No | The maximum number of scheduled actions per scalable target. |
| Step adjustments per step scaling policy | Each supported Region: 20 | Yes | The maximum number of step adjustments per step scaling policy. |

For more information, see Quotas for Application Auto Scaling in the *Application Auto Scaling User Guide*.

# AWS Application Discovery Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | discovery.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | discovery.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | discovery.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | discovery.ap-northeast-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | discovery.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | discovery.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | discovery.eu-west-2.amazonaws.com | HTTPS |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Active agents sending data to the service | Each supported Region: 1,000 | No | The maximum number of active agents sending data to the service. |
| Applications per account | Each supported Region: 1,000 | No | The maximum number of applications per account. If you reach this quota, and want to import new applications, you can delete existing applications in the Migration Hub console or by using the DeleteApplications API action. |
| Deletions of import records per day | Each supported Region: 25,000 | No | The maximum number of deletions of import records per day. Each day starts at 00:00 UTC. |
| Imported server records per account | Each supported Region: 25,000 | No | The maximum number of imported server records per account. |
| Imported servers per account | Each supported Region: 10,000 | Yes | The maximum number of imported servers per account. |
| Inactive agents heartbeating but not collecting data | Each supported Region: 10,000 | No | The maximum number of inactive agents heartbeating but not collecting data. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Servers per application | Each supported Region: 400 | No | The maximum number of servers that can be associated with a single application. |
| Tags per server | Each supported Region: 30 | No | The maximum number of tags for a single server. |

# AWS Transform MGN endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | mgn.us-east-2.amazonaws.com | HTTPS | |
| | | mgn-fips.us-east-2.api.aws | HTTPS | |
| | | mgn-fips.us-east-2.amazonaws.com | HTTPS | |
| | | mgn.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | mgn.us-east-1.amazonaws.com | HTTPS | |
| | | mgn-fips.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | mgn-fips.us-east-1.amazonaws.com | HTTPS |
| | | mgn.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | mgn.us-west-1.amazonaws.com | HTTPS |
| | | mgn-fips.us-west-1.api.aws | HTTPS |
| | | mgn-fips.us-west-1.amazonaws.com | HTTPS |
| | | mgn.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | mgn.us-west-2.amazonaws.com | HTTPS |
| | | mgn-fips.us-west-2.api.aws | HTTPS |
| | | mgn-fips.us-west-2.amazonaws.com | HTTPS |
| | | mgn.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | mgn.af-south-1.amazonaws.com | HTTPS |
| | | mgn.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | mgn.ap-east-1.amazonaws.com | HTTPS |
| | | mgn.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | mgn.ap-south-2.amazonaws.com | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | mgn.ap-southeast-3.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | mgn.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | mgn.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | mgn.ap-south-1.amazonaws.com<br>mgn.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | mgn.ap-southeast-6.amazonaws.com<br>mgn.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | mgn.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | mgn.ap-northeast-2.amazonaws.com<br>mgn.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | mgn.ap-southeast-1.amazonaws.com<br>mgn.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | mgn.ap-southeast-2.amazonaws.com<br>mgn.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | mgn.ap-east-2.amazonaws.com<br>mgn.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | mgn.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | mgn.ap-northeast-1.amazonaws.com<br>mgn.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | mgn.ca-central-1.amazonaws.com<br>mgn.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | mgn.ca-west-1.amazonaws.com<br>mgn.ca-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | mgn.eu-central-1.amazonaws.com<br>mgn.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | mgn.eu-west-1.amazonaws.com<br>mgn.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | mgn.eu-west-2.amazonaws.com<br>mgn.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | mgn.eu-south-1.amazonaws.com<br>mgn.eu-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | mgn.eu-west-3.amazonaws.com<br>mgn.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | mgn.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | mgn.eu-north-1.amazonaws.com<br>mgn.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | mgn.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | mgn.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | mgn.mx-central-1.amazonaws.com<br>mgn.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | mgn.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | mgn.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | mgn.sa-east-1.amazonaws.com<br>mgn.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | mgn.us-gov-east-1.amazonaws.com | HTTPS | |
| | | mgn-fips.us-gov-east-1.api.aws | HTTPS | |
| | | mgn-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | mgn.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | mgn.us-gov-west-1.amazonaws.com | HTTPS | |
| | | mgn-fips.us-gov-west-1.api.aws | HTTPS | |
| | | mgn-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | mgn.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Concurrent jobs in progress | Each supported Region: 20 | No | Concurrent jobs in progress |
| Concurrent jobs in progress | Each supported Region: 20 | No | Concurrent jobs in progress |
| Max Active Source Servers | Each supported Region: 150 | Yes | Max Active Source Servers |
| Max Active Source Servers | Each supported Region: 150 | Yes | Max Active Source Servers |
| Max Kubernetes applications per cluster | Each supported Region: 100 | No | Max Kubernetes applications per cluster |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max Kubernetes clusters per account | Each supported Region: 100 | No | Max Kubernetes clusters per account |
| Max NSX gateway policy rules per network migration definition | Each supported Region: 2,500 | Yes | Max NSX gateway policy rules per network migration definition |
| Max NSX gateway policy rules per network migration definition | Each supported Region: 2,500 | Yes | Max NSX gateway policy rules per network migration definition |
| Max NSX security policy rules per network migration definition | Each supported Region: 2,500 | Yes | Max NSX security policy rules per network migration definition |
| Max NSX security policy rules per network migration definition | Each supported Region: 2,500 | Yes | Max NSX security policy rules per network migration definition |
| Max Non-Archived Source Servers | Each supported Region: 4,000 | Yes | Max Non-Archived Source Servers |
| Max Non-Archived Source Servers | Each supported Region: 4,000 | Yes | Max Non-Archived Source Servers |
| Max Source Servers in a single Job | Each supported Region: 200 | No | Max Source Servers in a single Job |
| Max Source Servers in a single Job | Each supported Region: 200 | No | Max Source Servers in a single Job |
| Max Source Servers in all Jobs | Each supported Region: 200 | No | Max Source Servers in all Jobs |
| Max Source Servers in all Jobs | Each supported Region: 200 | No | Max Source Servers in all Jobs |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max Total Source Servers Per AWS Account | Each supported Region: 50,000 | No | Max Total Source Servers Per AWS Account |
| Max Total Source Servers Per AWS Account | Each supported Region: 50,000 | No | Max Total Source Servers Per AWS Account |
| Max actions per source server | Each supported Region: 100 | No | Max actions per source server |
| Max actions per source server | Each supported Region: 20 | No | Max actions per source server |
| Max actions per template | Each supported Region: 100 | No | Max actions per template |
| Max actions per template | Each supported Region: 20 | No | Max actions per template |
| Max active applications | Each supported Region: 200 | No | Max active applications |
| Max active applications | Each supported Region: 200 | No | Max active applications |
| Max active waves | Each supported Region: 200 | No | Max active waves |
| Max active waves | Each supported Region: 200 | No | Max active waves |
| Max applications per wave | Each supported Region: 1,000 | No | Max applications per wave |
| Max applications per wave | Each supported Region: 1,000 | No | Max applications per wave |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max archived applications | Each supported Region: 10,000 | No | Max archived applications |
| Max archived applications | Each supported Region: 10,000 | No | Max archived applications |
| Max archived waves | Each supported Region: 10,000 | No | Max archived waves |
| Max archived waves | Each supported Region: 10,000 | No | Max archived waves |
| Max concurrent Jobs per Source Server | Each supported Region: 1 | No | Max concurrent Jobs per Source Server |
| Max concurrent Jobs per Source Server | Each supported Region: 1 | No | Max concurrent Jobs per Source Server |
| Max conversions per Kubernetes cluster | Each supported Region: 5,000 | No | Max conversions per Kubernetes cluster |
| Max network migration definitions per account per Region | Each supported Region: 5 | Yes | Max network migration definitions per account per Region |
| Max network migration definitions per account per Region | Each supported Region: 5 | Yes | Max network migration definitions per account per Region |
| Max source servers per application | Each supported Region: 200 | No | Max source servers per application |
| Max source servers per application | Each supported Region: 200 | No | Max source servers per application |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Max target FSx file systems per account | Each supported Region: 5 | Yes | Max target FSx file systems per account |
| Max target FSx file systems per account | Each supported Region: 5 | Yes | Max target FSx file systems per account |

The following table lists additional information.

| Resource | Retention |
| --- | --- |
| Launch history | Saved for 10 years |
| Individual Job log | Saved for 185 days |

# Amazon AppStream 2.0 endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | appstream2.us-east-2.amazonaws.com | HTTPS | |
| | | appstream2.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | appstream2.us-east-1.amazonaws.com | HTTPS |
| | | appstream2-fips.us-east-1.api.aws | HTTPS |
| | | appstream2-fips.us-east-1.amazonaws.com | HTTPS |
| | | appstream2.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | appstream2.us-west-2.amazonaws.com | HTTPS |
| | | appstream2-fips.us-west-2.amazonaws.com | HTTPS |
| | | appstream2.us-west-2.api.aws | HTTPS |
| | | appstream2-fips.us-west-2.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | appstream2.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | appstream2.ap-south-1.amazonaws.com | HTTPS |
| | | appstream2.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | appstream2.ap-northeast-3.amazonaws.com | HTTPS |
| | | appstream2.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | appstream2.ap-northeast-2.amazonaws.com | HTTPS |
| | | appstream2.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | appstream2.ap-southeast-1.amazonaws.com | HTTPS |
| | | appstream2.ap-southeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | appstream2.ap-southeast-2.amazonaws.com | HTTPS |
| | | appstream2.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | appstream2.ap-northeast-1.amazonaws.com | HTTPS |
| | | appstream2.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | appstream2.ca-central-1.amazonaws.com | HTTPS |
| | | appstream2.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | appstream2.ca-west-1.amazonaws.com | HTTPS |
| | | appstream2.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | appstream2.eu-central-1.amazonaws.com | HTTPS |
| | | appstream2.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | appstream2.eu-west-1.amazonaws.com | HTTPS |
| | | appstream2.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | appstream2.eu-west-2.amazonaws.com | HTTPS |
| | | appstream2.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | appstream2.eu-south-1.amazonaws.com | HTTPS |
| | | appstream2.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | appstream2.eu-west-3.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | appstream2.eu-south-2.amazonaws.com | HTTPS |
| | | appstream2.eu-south-2.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | appstream2.eu-central-2.amazonaws.com | HTTPS |
| | | appstream2.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | appstream2.il-central-1.amazonaws.com | HTTPS |
| | | appstream2.il-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | appstream2.sa-east-1.amazonaws.com | HTTPS |
| | | appstream2.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | appstream2.us-gov-east-1.amazonaws.com | HTTPS |
| | | appstream2-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | appstream2-fips.us-gov-east-1.api.aws | HTTPS |
| | | appstream2.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | appstream2.us-gov-west-1.amazonaws.com | HTTPS |
| | | appstream2-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | appstream2.us-gov-west-1.api.aws | HTTPS |
| | | appstream2-fips.us-gov-west-1.api.aws | HTTPS |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Accelerated.g4dn.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g4dn.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g4dn.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g4dn.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.16xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g4dn.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g4dn.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g4dn.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g4dn.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g4dn.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g4dn.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g4dn.xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g4dn.xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g4dn.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g5.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.12xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Accelerated.g5.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g5.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g5.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.2xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g5.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g5.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g5.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.4xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g5.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g5.xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g5.xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g5.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Accelerated.g6.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.24xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g6.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.4xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Accelerated.g6.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6.xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6.xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6.xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g6e.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6e.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6e.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6e.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6e.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Accelerated.g6e.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6e.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6e.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6e.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6e.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.48xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g6e.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6e.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6e.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6e.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6e.xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g6e.xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6e.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6f.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6f.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6f.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6f.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.4xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g6f.large streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.large streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6f.large streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.large streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g6f.xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g6f.xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g6f.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g7.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g7.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g7.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g7.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g7.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g7.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.2xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g7.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g7.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g7.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.g7.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.g7.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Accelerated.g7.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.g7.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.gr6.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.gr6.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.gr6.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.gr6.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Accelerated.gr6.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.gr6.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.gr6.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.gr6.8xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Accelerated.gr6f.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of Accelerated.gr6f.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| Accelerated.gr6f.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of Accelerated.gr6f.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Active fleets | Each supported Region: 10 | Yes | The maximum number of active fleets that you can create in this account in the current Region. |
| AgentAccessMCP concurrent sessions | Each supported Region: 200 | Yes | The maximum number of concurrent active sessions for Agent Access MCP per account, in the current region. |
| App block builders | Each supported Region: 10 | Yes | The maximum number of app block builders that you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL compute-optimized 2xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL compute-optimized 2xlarge streaming instances (stream.compute.2xlarge) that you can use for fleets in this account in the current Region. |
| BYOL compute-optimized 2xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL compute-optimized 2xlarge streaming instances (stream.compute.2xlarge) that you can use for image builders in this account in the current Region. |
| BYOL compute-optimized 4xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL compute-optimized 4xlarge streaming instances (stream.compute.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL compute-optimized 4xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL compute-optimized 4xlarge streaming instances (stream.compute.4xlarge) that you can use for image builders in this account in the current Region. |
| BYOL compute-optimized 8xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL compute-optimized 8xlarge streaming instances (stream.compute.8xlarge) that you can use for fleets in this account in the current Region. |
| BYOL compute-optimized 8xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL compute-optimized 8xlarge streaming instances (stream.compute.8xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL compute-optimized large streaming instances for fleets | Each supported Region: 20 | Yes | The maximum number of BYOL compute-optimized large streaming instances (stream.compute.large) that you can use for fleets in this account in the current Region. |
| BYOL compute-optimized large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of BYOL compute-optimized large streaming instances (stream.compute.large) that you can use for image builders in this account in the current Region. |
| BYOL compute-optimized xlarge streaming instances for fleets | Each supported Region: 20 | Yes | The maximum number of BYOL compute-optimized xlarge streaming instances (stream.compute.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL compute-optimized xlarge streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of BYOL compute-optimized xlarge streaming instances (stream.compute.xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 12xlarge streaming instances (stream.graphics.g6.12xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 12xlarge image builder instances (stream.graphics.g6.12xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g6 16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 16xlarge streaming instances (stream.graphics.g6.16xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6 16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 16xlarge image builder instances (stream.graphics.g6.16xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6 24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 24xlarge streaming instances (stream.graphics.g6.24xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g6 24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 24xlarge image builder instances (stream.graphics.g6.24xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 2xlarge streaming instances (stream.graphics.g6.2xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 2xlarge image builder instances (stream.graphics.g6.2xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g6 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 4xlarge streaming instances (stream.graphics.g6.4xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 4xlarge image builder instances (stream.graphics.g6.4xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 8xlarge streaming instances (stream.graphics.g6.8xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g6 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 8xlarge image builder instances (stream.graphics.g6.8xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6 xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 xlarge streaming instances (stream.graphics.g6.xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6 xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6 xlarge image builder instances (stream.graphics.g6.xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| BYOL graphics g6f 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f 2xlarge streaming instances (stream.graphics.g6f.2xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6f 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f 2xlarge image builder instances (stream.graphics.g6f.2xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6f 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f 4xlarge streaming instances (stream.graphics.g6f.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g6f 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f 4xlarge image builder instances (stream.graphics.g6f.4xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6f large streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f large streaming instances (stream.graphics.g6f.large) that you can use for fleets in this account in the current Region. |
| BYOL graphics g6f large streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f large image builder instances (stream.graphics.g6f.large) that you can use for image builders in this account in the current Region. |
| BYOL graphics g6f xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f xlarge streaming instances (stream.graphics.g6f.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g6f xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g6f xlarge image builder instances (stream.graphics.g6f.xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g7 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 12xlarge streaming instances (stream.graphics.g7.12xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g7 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 12xlarge image builder instances (stream.graphics.g7.12xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g7 24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 24xlarge streaming instances (stream.graphics.g7.24xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g7 24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 24xlarge image builder instances (stream.graphics.g7.24xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g7 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 2xlarge streaming instances (stream.graphics.g7.2xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g7 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 2xlarge image builder instances (stream.graphics.g7.2xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g7 48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 48xlarge streaming instances (stream.graphics.g7.48xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g7 48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 48xlarge image builder instances (stream.graphics.g7.48xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g7 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 4xlarge streaming instances (stream.graphics.g7.4xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics g7 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 4xlarge image builder instances (stream.graphics.g7.4xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics g7 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 8xlarge streaming instances (stream.graphics.g7.8xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics g7 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics g7 8xlarge image builder instances (stream.graphics.g7.8xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics gr6 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics gr6 4xlarge streaming instances (stream.graphics.gr6.4xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics gr6 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics gr6 4xlarge image builder instances (stream.graphics.gr6.4xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics gr6 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics gr6 8xlarge streaming instances (stream.graphics.gr6.8xlarge) that you can use for fleets in this account in the current Region. |
| BYOL graphics gr6 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics gr6 8xlarge image builder instances (stream.graphics.gr6.8xlarge) that you can use for image builders in this account in the current Region. |
| BYOL graphics gr6f 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of BYOL graphics gr6f 4xlarge streaming instances (stream.graphics.gr6f.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL graphics gr6f 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of BYOL graphics gr6f 4xlarge image builder instances (stream.graphics.gr6f.4xlarge) that you can use for image builders in this account in the current Region. |
| BYOL memory-optimized 2xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL memory-optimized 2xlarge streaming instances (stream.memory.2xlarge) that you can use for fleets in this account in the current Region. |
| BYOL memory-optimized 2xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL memory-optimized 2xlarge streaming instances (stream.memory.2xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL memory-optimized 4xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL memory-optimized 4xlarge streaming instances (stream.memory.4xlarge) that you can use for fleets in this account in the current Region. |
| BYOL memory-optimized 4xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL memory-optimized 4xlarge streaming instances (stream.memory.4xlarge) that you can use for image builders in this account in the current Region. |
| BYOL memory-optimized 8xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL memory-optimized 8xlarge streaming instances (stream.memory.8xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL memory-optimized 8xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL memory-optimized 8xlarge streaming instances (stream.memory.8xlarge) that you can use for image builders in this account in the current Region. |
| BYOL memory-optimized large streaming instances for fleets | Each supported Region: 20 | Yes | The maximum number of BYOL memory-optimized large streaming instances (stream.memory.large) that you can use for fleets in this account in the current Region. |
| BYOL memory-optimized large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of BYOL memory-optimized large streaming instances (stream.memory.large) that you can use for image builders in this account in the current Region. |
| BYOL memory-optimized xlarge streaming instances for fleets | Each supported Region: 20 | Yes | The maximum number of BYOL memory-optimized xlarge streaming instances (stream.memory.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL memory-optimized xlarge streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of BYOL memory-optimized xlarge streaming instances (stream.memory.xlarge) that you can use for image builders in this account in the current Region. |
| BYOL standard 2xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of BYOL standard 2xlarge streaming instances (stream.standard.2xlarge) that you can use for fleets in this account in the current Region. |
| BYOL standard 2xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of BYOL standard 2xlarge streaming instances (stream.standard.2xlarge) that you can use for image builders in this account in the current Region. |
| BYOL standard large streaming instances for fleets | Each supported Region: 20 | Yes | The maximum number of BYOL standard large streaming instances (stream.standard.large) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BYOL standard large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of BYOL standard large streaming instances (stream.standard.large) that you can use for image builders in this account in the current Region. |
| BYOL standard medium streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of BYOL standard medium streaming instances (stream.standard.medium) that you can use for fleets in this account in the current Region. |
| BYOL standard medium streaming instances for image builders | Each supported Region: 10 | Yes | The maximum number of BYOL standard medium streaming instances (stream.standard.medium) that you can use for image builders in this account in the current Region. |
| BYOL standard xlarge streaming instances for fleets | Each supported Region: 20 | Yes | The maximum number of BYOL standard xlarge streaming instances (stream.standard.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| BYOL standard xlarge streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of BYOL standard xlarge streaming instances (stream.standard.xlarge) that you can use for image builders in this account in the current Region. |
| Compute-optimized 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of compute-optimized 2xlarge streaming instances (stream.compute.2xlarge) that you can use for fleets in this account in the current Region. |
| Compute-optimized 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of compute-optimized 2xlarge streaming instances (stream.compute.2xlarge) that you can use for image builders in this account in the current Region. |
| Compute-optimized 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of compute 4xlarge streaming instances (stream.compute.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Compute-optimized 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of compute-optimized 4xlarge streaming instances (stream.compute.4xlarge) that you can use for image builders in this account in the current Region. |
| Compute-optimized 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of compute-optimized 8xlarge streaming instances (stream.compute.8xlarge) that you can use for fleets in this account in the current Region. |
| Compute-optimized 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of compute-optimized 8xlarge streaming instances (stream.compute.8xlarge) that you can use for image builders in this account in the current Region. |
| Compute-optimized large streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of compute-optimized large streaming instances (stream.compute.large) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Compute-optimized large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of compute-optimized large streaming instances (stream.compute.large) that you can use for image builders in this account in the current Region. |
| Compute-optimized xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of compute-optimized xlarge streaming instances (stream.compute.xlarge) that you can use for fleets in this account in the current Region. |
| Compute-optimized xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of compute-optimized xlarge streaming instances (stream.compute.xlarge) that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6a.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6a.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6a.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6a.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.32xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6a.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6a.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.4xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6a.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6a.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6a.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c6a.large streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6a.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c6a.large streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6a.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c6a.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6a.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c6a.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6i.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6i.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6i.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6i.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.32xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6i.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6i.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6i.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c6i.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c6i.large streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c6i.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c6i.large streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c6i.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c6i.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c6i.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c6i.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7a.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.16xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7a.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7a.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.2xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7a.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.32xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7a.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.48xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7a.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7a.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7a.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7a.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c7a.large streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7a.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c7a.large streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7a.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c7a.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7a.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c7a.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i-flex.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.16xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i-flex.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i-flex.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.4xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i-flex.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i-flex.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i-flex.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c7i-flex.large streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i-flex.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c7i-flex.large streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i-flex.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c7i-flex.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i-flex.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c7i-flex.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.48xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of ComputeOptimized.c7i.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c7i.large streaming instances that you can use for fleets in this account in the current Region. |
| ComputeOptimized.c7i.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of ComputeOptimized.c7i.large streaming instances that you can use for image builders in this account in the current Region. |
| ComputeOptimized.c7i.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c7i.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ComputeOptimized.c7i.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of ComputeOptimized.c7i.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Concurrent image copies per destination Region | Each supported Region: 2 | Yes | The maximum number of concurrent image copies that you can have in this account in the current Region. |
| Concurrent image updates | Each supported Region: 5 | Yes | The maximum number of concurrent image updates that you can have in this account in the current Region. |
| Fleets | Each supported Region: 10 | Yes | The maximum number of fleets that you can create in this account in the current Region. |
| GeneralPurpose.m6a.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6a.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6a.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6a.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6a.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.32xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6a.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6a.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.4xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| GeneralPurpose.m6a.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6a.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6a.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.m6a.large streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6a.large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m6a.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6a.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of GeneralPurpose.m6a.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6a.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m6a.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6i.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6i.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| GeneralPurpose.m6i.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6i.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.32xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6i.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6i.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6i.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m6i.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.m6i.large streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m6i.large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m6i.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m6i.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of GeneralPurpose.m6i.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m6i.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m6i.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7a.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.16xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7a.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7a.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.2xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7a.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.32xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7a.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.48xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7a.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7a.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7a.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7a.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.m7a.large streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7a.large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m7a.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7a.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of GeneralPurpose.m7a.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7a.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m7a.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i-flex.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.16xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i-flex.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i-flex.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.4xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i-flex.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i-flex.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i-flex.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.m7i-flex.large streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i-flex.large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m7i-flex.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i-flex.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of GeneralPurpose.m7i-flex.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i-flex.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m7i-flex.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.12xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.24xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| GeneralPurpose.m7i.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.48xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.8xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of GeneralPurpose.m7i.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.m7i.large streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.m7i.large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m7i.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.m7i.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of GeneralPurpose.m7i.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.m7i.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.m7i.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3.2xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of GeneralPurpose.t3.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.t3.2xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.t3.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3.large streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of GeneralPurpose.t3.large streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.t3.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.t3.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3.medium streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of GeneralPurpose.t3.medium streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.t3.medium streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.t3.medium streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3.small streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of GeneralPurpose.t3.small streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| GeneralPurpose.t3.small streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.t3.small streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3.xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of GeneralPurpose.t3.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.t3.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.t3.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3a.2xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of GeneralPurpose.t3a .2xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.t3a.2xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.t3a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3a.large streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of GeneralPurpose.t3a.large streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.t3a.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.t3a.large streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3a.medium streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of GeneralPurpose.t3a.medium streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.t3a.medium streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.t3a.medium streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3a.small streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of GeneralPurpose.t3a.small streaming instances that you can use for fleets in this account in the current Region. |
| GeneralPurpose.t3a.small streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of GeneralPurpose.t3a.small streaming instances that you can use for image builders in this account in the current Region. |
| GeneralPurpose.t3a.xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of GeneralPurpose.t3a.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GeneralPurpose.t3a.xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of GeneralPurpose.t3a.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Graphics G4DN 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 12xlarge instances (stream.graphics.g4dn.12xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G4DN 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 12xlarge instances (stream.graphics.g4dn.12xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G4DN 16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 16xlarge instances (stream.graphics.g4dn.16xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G4DN 16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 16xlarge instances (stream.graphics.g4dn.16xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G4DN 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 2xlarge instances (stream.graphics.g4dn.2xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G4DN 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 2xlarge instances (stream.graphics.g4dn.2xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G4DN 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 4xlarge instances (stream.graphics.g4dn.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G4DN 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 4xlarge instances (stream.graphics.g4dn.4xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G4DN 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 8xlarge instances (stream.graphics.g4dn.8xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G4DN 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G4DN 8xlarge instances (stream.graphics.g4dn.8xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G4DN xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G4DN xlarge instances (stream.graphics.g4dn.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Graphics G4DN xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G4DN xlarge instances (stream.graphics.g4dn.xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 12xlarge instances (stream.graphics.g5.12xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G5 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 12xlarge instances (stream.graphics.g5.12xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 16xlarge instances (stream.graphics.g5.16xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Graphics G5 16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 16xlarge instances (stream.graphics.g5.16xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 24xlarge instances (stream.graphics.g5.24xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G5 24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 24xlarge instances (stream.graphics.g5.24xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 2xlarge instances (stream.graphics.g5.2xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G5 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 2xlarge instances (stream.graphics.g5.2xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 4xlarge instances (stream.graphics.g5.4xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G5 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 4xlarge instances (stream.graphics.g5.4xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 8xlarge instances (stream.graphics.g5.8xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G5 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 8xlarge instances (stream.graphics.g5.8xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G5 xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G5 xlarge instances (stream.graphics.g5.xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G5 xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G5 xlarge instances (stream.graphics.g5.xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 12xlarge instances (stream.graphics.g6.12xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G6 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 12xlarge instances (stream.graphics.g6.12xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 16xlarge instances (stream.graphics.g6.16xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G6 16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 16xlarge instances (stream.graphics.g6.16xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 24xlarge instances (stream.graphics.g6.24xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G6 24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 24xlarge instances (stream.graphics.g6.24xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 2xlarge instances (stream.graphics.g6.2xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G6 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 2xlarge instances (stream.graphics.g6.2xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 4xlarge instances (stream.graphics.g6.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G6 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 4xlarge instances (stream.graphics.g6.4xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 8xlarge instances (stream.graphics.g6.8xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G6 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 8xlarge instances (stream.graphics.g6.8xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6 xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6 xlarge instances (stream.graphics.g6.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Graphics G6 xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6 xlarge instances (stream.graphics.g6.xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6F 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6F 2xlarge instances (stream.graphics.g6f.2xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G6F 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6F 2xlarge instances (stream.graphics.g6f.2xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6F 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6F 4xlarge instances (stream.graphics.g6f.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Graphics G6F 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6F 4xlarge instances (stream.graphics.g6f.4xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G6F large streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6F large instances (stream.graphics.g6f.large) that you can use for fleets in this account in the current Region. |
| Graphics G6F large streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6F large instances (stream.graphics.g6f.large) that you can use for image builders in this account in the current Region. |
| Graphics G6F xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G6F xlarge instances (stream.graphics.g6f.xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G6F xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G6F xlarge instances (stream.graphics.g6f.xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G7 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G7 12xlarge instances (stream.graphics.g7.12xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G7 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G7 12xlarge instances (stream.graphics.g7.12xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G7 24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G7 24xlarge instances (stream.graphics.g7.24xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G7 24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G7 24xlarge instances (stream.graphics.g7.24xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G7 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G7 2xlarge instances (stream.graphics.g7.2xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G7 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G7 2xlarge instances (stream.graphics.g7.2xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G7 48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G7 48xlarge instances (stream.graphics.g7.48xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G7 48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G7 48xlarge instances (stream.graphics.g7.48xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G7 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G7 4xlarge instances (stream.graphics.g7.4xlarge) that you can use for fleets in this account in the current Region. |
| Graphics G7 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G7 4xlarge instances (stream.graphics.g7.4xlarge) that you can use for image builders in this account in the current Region. |
| Graphics G7 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics G7 8xlarge instances (stream.graphics.g7.8xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Graphics G7 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics G7 8xlarge instances (stream.graphics.g7.8xlarge) that you can use for image builders in this account in the current Region. |
| Graphics GR6 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics GR6 4xlarge instances (stream.graphics.gr6.4xlarge) that you can use for fleets in this account in the current Region. |
| Graphics GR6 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics GR6 4xlarge instances (stream.graphics.gr6.4xlarge) that you can use for image builders in this account in the current Region. |
| Graphics GR6 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics GR6 8xlarge instances (stream.graphics.gr6.8xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Graphics GR6 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics GR6 8xlarge instances (stream.graphics.gr6.8xlarge) that you can use for image builders in this account in the current Region. |
| Graphics GR6F 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of graphics GR6F 4xlarge instances (stream.graphics.gr6f.4xlarge) that you can use for fleets in this account in the current Region. |
| Graphics GR6F 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of graphics GR6F 4xlarge instances (stream.graphics.gr6f.4xlarge) that you can use for image builders in this account in the current Region. |
| Image builders | Each supported Region: 10 | Yes | The maximum number of image builders that you can create in this account in the current Region. |
| Image sharing limit | Each supported Region: 100 | Yes | The maximum number of AWS accounts one image can be shared with in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max app block builders with Windows Server 2019 platform and stream.standard.2xlarge instance type | Each supported Region: 1 | Yes | The maximum number of app block builders with Windows Server 2019 platform and stream.standard.2xlarge instance type you can use in this account in the current Region. |
| Max app block builders with Windows Server 2019 platform and stream.standard.large instance type | Each supported Region: 1 | Yes | The maximum number of app block builders with Windows Server 2019 platform and stream.standard.large instance type you can use in this account in the current Region. |
| Max app block builders with Windows Server 2019 platform and stream.standard.medium instance type | Each supported Region: 1 | Yes | The maximum number of app block builders with Windows Server 2019 platform and stream.standard.medium instance type you can use in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max app block builders with Windows Server 2019 platform and stream.standard.small instance type | Each supported Region: 1 | Yes | The maximum number of app block builders with Windows Server 2019 platform and stream.standard.small instance type you can use in this account in the current Region. |
| Max app block builders with Windows Server 2019 platform and stream.standard.xlarge instance type | Each supported Region: 1 | Yes | The maximum number of app block builders with Windows Server 2019 platform and stream.standard.xlarge instance type you can use in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Ubuntu Pro 2404 platform and stream.standard.2xlarge instance type | Each supported Region: 2 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Ubuntu Pro 2404 platform and stream.standard.2xlarge instance type in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max concurrent sessions for Elastic fleets with Ubuntu Pro 2404 platform and stream.standard.large instance type | Each supported Region: 5 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Ubuntu Pro 2404 platform and stream.standard.large instance type in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Ubuntu Pro 2404 platform and stream.standard.medium instance type | Each supported Region: 10 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Ubuntu Pro 2404 platform and stream.standard.medium instance type in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Ubuntu Pro 2404 platform and stream.standard.small instance type | Each supported Region: 10 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Ubuntu Pro 2404 platform and stream.standard.small instance type in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max concurrent sessions for Elastic fleets with Ubuntu Pro 2404 platform and stream.standard.xlarge instance type | Each supported Region: 2 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Ubuntu Pro 2404 platform and stream.standard.xlarge instance type in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Windows Server 2019 platform and stream.standard.2xlarge instance type | Each supported Region: 2 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Windows Server 2019 platform and stream.standard.2xlarge instance type in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Windows Server 2019 platform and stream.standard.large instance type | Each supported Region: 5 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Windows Server 2019 platform and stream.standard.large instance type in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max concurrent sessions for Elastic fleets with Windows Server 2019 platform and stream.standard.medium instance type | Each supported Region: 10 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Windows Server 2019 platform and stream.standard.medium instance type in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Windows Server 2019 platform and stream.standard.small instance type | Each supported Region: 10 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Windows Server 2019 platform and stream.standard.small instance type in this account in the current Region. |
| Max concurrent sessions for Elastic fleets with Windows Server 2019 platform and stream.standard.xlarge instance type | Each supported Region: 2 | Yes | The maximum number of concurrent sessions that you can use for Elastic fleets Windows Server 2019 platform and stream.standard.xlarge instance type in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized 2xlarge streaming instances (stream.memory.2xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized 2xlarge streaming instances (stream.memory.2xlarge) that you can use for image builders in this account in the current Region. |
| Memory-optimized 4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized 4xlarge streaming instances (stream.memory.4xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized 4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized 4xlarge streaming instances (stream.memory.4xlarge) that you can use for image builders in this account in the current Region. |
| Memory-optimized 8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized 8xlarge streaming instances (stream.memory.8xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized 8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized 8xlarge streaming instances (stream.memory.8xlarge) that you can use for image builders in this account in the current Region. |
| Memory-optimized large streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of memory-optimized large streaming instances (stream.memory.large) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of memory-optimized large streaming instances (stream.memory.large) that you can use for image builders in this account in the current Region. |
| Memory-optimized xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of memory-optimized xlarge streaming instances (stream.memory.xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of memory-optimized xlarge streaming instances (stream.memory.xlarge) that you can use for image builders in this account in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized z1d 12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 12xlarge streaming instances (stream.memory.z1d.12xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized z1d 12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 12xlarge streaming instances (stream.memory.z1d.12xlarge) that you can use for image builders in this account in the current Region. |
| Memory-optimized z1d 2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 2xlarge streaming instances (stream.memory.z1d.2xlarge) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized z1d 2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 2xlarge streaming instances (stream.memory.z1d.2xlarge) that you can use for image builders in this account in the current Region. |
| Memory-optimized z1d 3xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 3xlarge streaming instances (stream.memory.z1d.3xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized z1d 3xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 3xlarge streaming instances (stream.memory.z1d.3xlarge) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized z1d 6xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 6xlarge streaming instances (stream.memory.z1d.6xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized z1d 6xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of memory-optimized z1d 6xlarge streaming instances (stream.memory.z1d.6xlarge) that you can use for image builders in this account in the current Region. |
| Memory-optimized z1d large streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of memory-optimized z1d large streaming instances (stream.memory.z1d.large) that you can use for fleets in this account in the current Region. |
| Memory-optimized z1d large streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of memory-optimized z1d large streaming instances (stream.memory.z1d.large) that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Memory-optimized z1d xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of memory-optimized z1d xlarge streaming instances (stream.memory.z1d.xlarge) that you can use for fleets in this account in the current Region. |
| Memory-optimized z1d xlarge streaming instances for image builders | Each supported Region: 3 | Yes | The maximum number of memory-optimized z1d xlarge streaming instances (stream.memory.z1d.xlarge) that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6a.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.12xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6a.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6a.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.24xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6a.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6a.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.32xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.32xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6a.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6a.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.4xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6a.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6a.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6a.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r6a.large streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r6a.large streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6a.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r6a.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6a.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r6a.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6i.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.12xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6i.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.16xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6i.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.24xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6i.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.2xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6i.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.32xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.32xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6i.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6i.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r6i.8xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r6i.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r6i.large streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r6i.large streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r6i.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r6i.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r6i.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r6i.xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7a.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7a.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.16xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7a.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7a.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.2xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7a.32xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.32xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.32xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.32xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7a.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.48xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7a.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.4xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7a.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7a.8xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7a.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r7a.large streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r7a.large streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7a.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r7a.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7a.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r7a.xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7i.12xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.12xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.12xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.12xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7i.16xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.16xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.16xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.16xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7i.24xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.24xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.24xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.24xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7i.2xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.2xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.2xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.2xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7i.48xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.48xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.48xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.48xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7i.4xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.4xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.4xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.4xlarge streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7i.8xlarge streaming instances for fleets | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.8xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.8xlarge streaming instances for image builders | Each supported Region: 0 | Yes | The maximum number of MemoryOptimized.r7i.8xlarge streaming instances that you can use for image builders in this account in the current Region. |
| MemoryOptimized.r7i.large streaming instances for fleets | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r7i.large streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of MemoryOptimized.r7i.large streaming instances that you can use for image builders in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| MemoryOptimized.r7i.xlarge streaming instances for fleets | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r7i.xlarge streaming instances that you can use for fleets in this account in the current Region. |
| MemoryOptimized.r7i.xlarge streaming instances for image builders | Each supported Region: 2 | Yes | The maximum number of MemoryOptimized.r7i.xlarge streaming instances that you can use for image builders in this account in the current Region. |
| Private images | Each supported Region: 10 | Yes | The maximum number of private images that you can create in this account in the current Region. |
| Stacks | Each supported Region: 10 | Yes | The maximum number of stacks that you can create in this account in the current Region. |
| Standard 2xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of stream.standard.2xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Standard 2xlarge streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of standard 2xlarge streaming instances (stream.standard.2xlarge) that you can use for image builders in this account in the current Region. |
| Standard large streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of standard large streaming instances (stream.standard.large) that you can use for fleets in this account in the current Region. |
| Standard large streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of standard large streaming instances (stream.standard.large) that you can use for image builders in this account in the current Region. |
| Standard medium streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of standard medium streaming instances (stream.standard.medium) that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Standard medium streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of standard medium streaming instances (stream.standard.medium) that you can use for image builders in this account in the current Region. |
| Standard small streaming instances for fleets | Each supported Region: 50 | Yes | The maximum number of standard small streaming instances (stream.standard.small) that you can use for fleets in this account in the current Region. |
| Standard small streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of standard small streaming instances (stream.standard.small) that you can use for image builders in this account in the current Region. |
| Standard xlarge streaming instances for fleets | Each supported Region: 10 | Yes | The maximum number of stream.standard.xlarge streaming instances that you can use for fleets in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Standard xlarge streaming instances for image builders | Each supported Region: 5 | Yes | The maximum number of standard xlarge streaming instances (stream.standard.xlarge) that you can use for image builders in this account in the current Region. |
| Throttle rate limit for AgentAccessMCP | Each supported Region: 10 | No | The maximum number of AgentAccessMCP requests you can make per second per account, in the current region |
| Users in the user pool | Each supported Region: 50 | Yes | The maximum number of users that you can create in the user pool in this account in the current Region. |

*For fleets that have **Default Internet Access** enabled, the quota is 100 fleet instances. If your deployment must support more than 100 concurrent users, use a NAT gateway configuration instead.

## AWS App Studio endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | appstudio.us-west-2.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | appstudio.eu-west-1.api.aws | HTTPS | |

## Service quotas

For information on quotas, see Quotas for AWS App Studio in the *AWS App Studio User Guide*.

# AWS AppSync endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### AWS AppSync Control plane

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | appsync.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | appsync.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | appsync.us-east-1.amazonaws.com<br>appsync.us-east-1.api.aws | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | appsync.us-west-1.amazonaws.com<br>appsync.us-west-1.api.aws | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | appsync.us-west-2.amazonaws.com<br>appsync.us-west-2.api.aws | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | appsync.af-south-1.amazonaws.com<br>appsync.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | appsync.ap-east-1.amazonaws.com<br>appsync.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | appsync.ap-south-2.amazonaws.com<br>appsync.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | appsync.ap-southeast-3.amazonaws.com<br>appsync.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | appsync.ap-southeast-5.amazonaws.com<br>appsync.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | appsync.ap-southeast-4.amazonaws.com<br>appsync.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | appsync.ap-south-1.amazonaws.com<br>appsync.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | appsync.ap-northeast-3.amazonaws.com<br>appsync.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | appsync.ap-northeast-2.amazonaws.com<br>appsync.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | appsync.ap-southeast-1.amazonaws.com<br>appsync.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | appsync.ap-southeast-2.amazonaws.com<br>appsync.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | appsync.ap-southeast-7.amazonaws.com<br>appsync.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | appsync.ap-northeast-1.amazonaws.com<br>appsync.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | appsync.ca-central-1.amazonaws.com<br>appsync.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | appsync.ca-west-1.amazonaws.com<br>appsync.ca-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | appsync.eu-central-1.amazonaws.com<br>appsync.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | appsync.eu-west-1.amazonaws.com<br>appsync.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | appsync.eu-west-2.amazonaws.com<br>appsync.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | appsync.eu-south-1.amazonaws.com<br>appsync.eu-south-1.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | appsync.eu-west-3.amazonaws.com<br>appsync.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | appsync.eu-south-2.amazonaws.com<br>appsync.eu-south-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | appsync.eu-north-1.amazonaws.com<br>appsync.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | appsync.eu-central-2.amazonaws.com<br>appsync.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | appsync.il-central-1.amazonaws.com<br>appsync.il-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | appsync.me-south-1.amazonaws.com<br>appsync.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | appsync.me-central-1.amazonaws.com<br>appsync.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | appsync.sa-east-1.amazonaws.com<br>appsync.sa-east-1.api.aws | HTTPS<br>HTTPS |

## AWS AppSync Gateway

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | appsync-api.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | appsync-api.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | appsync-api.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | appsync-api.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | appsync-api.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | appsync-api.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | appsync-api.ap-northeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | appsync-api.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | appsync-api.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | appsync-api.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | appsync-api.ap-south-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | appsync-api.ap-southeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | appsync-api.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | appsync-api.ap-southeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | appsync-api.ap-southeast-4.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | appsync-api.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | appsync-api.ap-southeast-7.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | appsync-api.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | appsync-api.ca-west-1.amazonaws.com | HTTPS |
| China (Beijing) | cn-north-1 | appsync-api.cn-north-1.amazonaws.com.cn | HTTPS |
| China (Ningxia) | cn-northwest-1 | appsync-api.cn-northwest-1.amazonaws.com.cn | HTTPS |
| Europe (Frankfurt) | eu-central-1 | appsync-api.eu-central-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | appsync-api.eu-central-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | appsync-api.eu-north-1.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | appsync-api.eu-south-1.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | appsync-api.eu-south-2.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | appsync-api.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | appsync-api.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | appsync-api.eu-west-3.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | appsync-api.il-central-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | appsync-api.me-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | appsync-api.me-south-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | appsync-api.sa-east-1.amazonaws.com | HTTPS |

## AWS AppSync Real-time Subscription

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | appsync-realtime-api.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | appsync-realtime-api.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | appsync-realtime-api.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | appsync-realtime-api.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | appsync-realtime-api.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | appsync-realtime-api.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | appsync-realtime-api.ap-northeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | appsync-realtime-api.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | appsync-realtime-api.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | appsync-realtime-api.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | appsync-realtime-api.ap-south-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | appsync-realtime-api.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | appsync-realtime-api.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | appsync-realtime-api.ap-southeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | appsync-realtime-api.ap-southeast-4.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | appsync-realtime-api.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | appsync-realtime-api.ap-southeast-7.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | appsync-realtime-api.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | appsync-realtime-api.ca-west-1.amazonaws.com | HTTPS |
| China (Beijing) | cn-north-1 | appsync-realtime-api.cn-north-1.amazonaws.com.cn | HTTPS |
| China (Ningxia) | cn-northwest-1 | appsync-realtime-api.cn-northwest-1.amazonaws.com.cn | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | appsync-realtime-api.eu-central-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | appsync-realtime-api.eu-central-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | appsync-realtime-api.eu-north-1.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | appsync-realtime-api.eu-south-1.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | appsync-realtime-api.eu-south-2.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | appsync-realtime-api.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | appsync-realtime-api.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | appsync-realtime-api.eu-west-3.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | appsync-realtime-api.il-central-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | appsync-realtime-api.me-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | appsync-realtime-api.me-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| South America (São Paulo) | sa-east-1 | appsync-realtime-api.sa-east-1.amazonaws.com | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| All APIs - API keys per API | Each supported Region: 50 | No | Maximum number of API keys per API |
| All APIs - Authentication providers per API | Each supported Region: 50 | No | The maximum number of authentication providers per API |
| All APIs - Functions per pipeline resolver or handler | Each supported Region: 10 | No | Maximum number of functions per pipeline resolver or handler |
| All APIs - Handler, resolver, and function code size | Each supported Region: 32 Kilobytes | No | Maximum code size of handlers, resolvers, and functions running on APPSYNC_JS runtime. |
| All APIs - Iterations per loops in resolvers, functions, and handlers | Each supported Region: 1,000 | No | Maximum number of iterations in #foreach. ..#end loops |
| All APIs - Number of custom domain names | Each supported Region: 50 | Yes | The maximum number of custom domain names per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| All APIs - Request execution time | Each supported Region: 30 Seconds | No | Maximum request execution time |
| All APIs - Resolvers, functions, and handlers response size | Each supported Region: 5 Megabytes | No | Maximum size of the evaluated resolver, function, or handler template |
| All APIs - Subscription payload size | Each supported Region: 240 Kilobytes | No | Maximum size of the message received via subscriptions |
| All APIs - Subscriptions per client connection | Each supported Region: 200 | Yes | Subscriptions code statements per client in a single WebSocket connection. Each statement allows the client to receive multiple messages related to the specific subscription. |
| Event APIs - APIs per region | Each supported Region: 50 | Yes | The maximum number of AppSync Event APIs per region |
| Event APIs - Batch size per publish request | Each supported Region: 5 | No | The maximum number of events that can be batched in one publish operation. |
| Event APIs - Channel namespaces per API | Each supported Region: 50 | Yes | The maximum number of channel namespaces on an Event API |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Event APIs - Number of characters allowed in a channel segment | Each supported Region: 50 | No | The maximum character length of a channel segment |
| Event APIs - Number of segments allowed in a channel | Each supported Region: 5 | No | The maximum number of segments in a channel |
| Event APIs - Publish payload size | Each supported Region: 1.2 Megabytes | No | The maximum size of the entire payload sent in a publish operation |
| Event APIs - Rate of connections per API | Each supported Region: 2,000 per second | Yes | The maximum number of Event API connect requests per second. |
| Event APIs - Rate of inbound events per API | Each supported Region: 10,000 per second | Yes | The maximum number of inbound events per second. Each published event generates a single inbound event. |
| Event APIs - Rate of outbound messages per API | Each supported Region: 1,000,000 per second | Yes | The maximum number of outbound metered events per second. One metered event equals 5 kB of delivered event. |
| Event APIs - Rate of publish requests per client connection | Each supported Region: 25 per second | No | The maximum number of publish requests per second per WebSocket connection. WebSocket publish requests are also evaluated against the rate of request tokens quota. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Event APIs - Rate of request tokens | Each supported Region: 2,000 per second | Yes | The maximum number of request tokens per second for Event APIs in this account in the current Region. AWS AppSync allocates tokens to publish requests based on the amount of resources (processing time and memory) that they consume and the number of 5kB messages published. |
| GraphQL APIs - APIs per region | Each supported Region: 50 | Yes | Maximum number of GraphQL APIs per region |
| GraphQL APIs - Max Batch Size per request | Each supported Region: 2,000 | No | The maximum length of the resolver request list that will be sent to a single Lambda function for a BatchInvoke operation. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GraphQL APIs - Maximum query length | Each supported Region: 15,000 | No | Maximum length of a GraphQL query, measured in tokens (15,000). Tokens are individual parsed elements of the query string. Queries exceeding this limit will fail with the error message Maximum char length exceeded. |
| GraphQL APIs - Number of caching keys | Each supported Region: 25 | No | The maximum number of caching keys |
| GraphQL APIs - Rate of connections per API | Each supported Region: 2,000 per second | Yes | The maximum number of GraphQL API connect requests per second. |
| GraphQL APIs - Rate of inbound messages per API | Each supported Region: 10,000 per second | Yes | The maximum number of inbound messages per second. Each subscription type invoked by a mutation generates one inbound message. |
| GraphQL APIs - Rate of outbound messages per API | Each supported Region: 1,000,000 per second | Yes | The maximum number of outbound messages per second. One outbound message is equal to 5 KB of payload data delivered. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GraphQL APIs - Rate of request tokens | us-east-1: 10,000 per second<br><br>us-east-2: 10,000 per second<br><br>us-west-2: 10,000 per second<br><br>ap-northeast-1: 10,000 per second<br><br>ap-northeast-2: 10,000 per second<br><br>ap-south-1: 10,000 per second<br><br>ap-southeast-1: 10,000 per second<br><br>ap-southeast-2: 10,000 per second<br><br>eu-central-1: 10,000 per second<br><br>eu-west-1: 10,000 per second | Yes | The maximum number of request tokens per second for GraphQL APIs in this account in the current Region. AWS AppSync allocates tokens to mutation and query requests based on the amount of resources (processing time and memory) that they consume. For more details on tokens, see the Monitoring section in the AWS AppSync documentation. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | eu-west-2: 10,000 per second<br><br>Each of the other supported Regions: 5,000 per second | | |
| GraphQL APIs - Request mapping template size | Each supported Region: 64 Kilobytes | No | Maximum size of the request mapping template |
| GraphQL APIs - Resolvers executed in a single request | Each supported Region: 10,000 | No | The maximum number of resolvers that can be executed in a single request |
| GraphQL APIs - Response mapping template size | Each supported Region: 64 Kilobytes | No | Maximum size of the response mapping template |
| GraphQL APIs - Schema document size | Each supported Region: 1 Megabytes | No | Maximum size of the schema document |
| GraphQL APIs - Source API associations per Merged API | Each supported Region: 10 | Yes | The maximum number of Source API associations per Merged API. |
| Graphql APIs - Rate of subscription invalidation requests | Each supported Region: 100 per second | No | The maximum number of invalidation requests per second per account per region |

**Rate of request tokens** is the maximum number of request tokens per second in this account in the current Region. AWS AppSync allocates tokens to mutation and query requests based on the amount of resources (processing time and memory) that they consume. For more details on tokens, see Using token counts to optimize your requests in the *AWS AppSync developer guide*.

# Amazon Athena endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | athena.us-east-2.amazonaws.com | HTTPS |
| | | athena-fips.us-east-2.api.aws | HTTPS |
| | | athena-fips.us-east-2.amazonaws.com | HTTPS |
| | | athena.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | athena.us-east-1.amazonaws.com | HTTPS |
| | | athena-fips.us-east-1.api.aws | HTTPS |
| | | athena-fips.us-east-1.amazonaws.com | HTTPS |
| | | athena.us-east-1.api.aws | HTTPS |
| US West (N. | us-west-1 | athena.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| California) | | athena-fips.us-west-1.api.aws | HTTPS | |
| | | athena-fips.us-west-1.amazonaws.com | HTTPS | |
| | | athena.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | athena.us-west-2.amazonaws.com | HTTPS | |
| | | athena-fips.us-west-2.api.aws | HTTPS | |
| | | athena-fips.us-west-2.amazonaws.com | HTTPS | |
| | | athena.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | athena.af-south-1.amazonaws.com | HTTPS | |
| | | athena.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | athena.ap-east-1.amazonaws.com | HTTPS | |
| | | athena.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | athena.ap-south-2.amazonaws.com | HTTPS | |
| | | athena.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | athena.ap-southeast-3.amazonaws.com | HTTPS | |
| | | athena.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | athena.ap-southeast-5.amazonaws.com | HTTPS | |
| | | athena.ap-southeast-5.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | athena.ap-southeast-4.amazonaws.com<br>athena.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | athena.ap-south-1.amazonaws.com<br>athena.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | athena.ap-southeast-6.amazonaws.com<br>athena.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | athena.ap-northeast-3.amazonaws.com<br>athena.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | athena.ap-northeast-2.amazonaws.com<br>athena.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | athena.ap-southeast-1.amazonaws.com<br>athena.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | athena.ap-southeast-2.amazonaws.com<br>athena.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | athena.ap-east-2.amazonaws.com<br>athena.ap-east-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | athena.ap-southeast-7.amazonaws.com | HTTPS |
| | | athena.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | athena.ap-northeast-1.amazonaws.com | HTTPS |
| | | athena.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | athena.ca-central-1.amazonaws.com | HTTPS |
| | | athena-fips.ca-central-1.api.aws | HTTPS |
| | | athena-fips.ca-central-1.amazonaws.com | HTTPS |
| | | athena.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | athena.ca-west-1.amazonaws.com | HTTPS |
| | | athena-fips.ca-west-1.api.aws | HTTPS |
| | | athena-fips.ca-west-1.amazonaws.com | HTTPS |
| | | athena.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | athena.eu-central-1.amazonaws.com | HTTPS |
| | | athena.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | athena.eu-west-1.amazonaws.com | HTTPS |
| | | athena.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | athena.eu-west-2.amazonaws.com | HTTPS |
| | | athena.eu-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | athena.eu-south-1.amazonaws.com | HTTPS | |
| | | athena.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | athena.eu-west-3.amazonaws.com | HTTPS | |
| | | athena.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | athena.eu-south-2.amazonaws.com | HTTPS | |
| | | athena.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | athena.eu-north-1.amazonaws.com | HTTPS | |
| | | athena.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | athena.eu-central-2.amazonaws.com | HTTPS | |
| | | athena.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | athena.il-central-1.amazonaws.com | HTTPS | |
| | | athena.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | athena.mx-central-1.amazonaws.com | HTTPS | |
| | | athena.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | athena.me-south-1.amazonaws.com | HTTPS | |
| | | athena.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | athena.me-central-1.amazonaws.com | HTTPS | |
| | | athena.me-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| South America (São Paulo) | sa-east-1 | athena.sa-east-1.amazonaws.com<br>athena.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | athena.us-gov-east-1.amazonaws.com<br>athena-fips.us-gov-east-1.api.aws<br>athena-fips.us-gov-east-1.amazonaws.com<br>athena.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | athena.us-gov-west-1.amazonaws.com<br>athena-fips.us-gov-west-1.api.aws<br>athena-fips.us-gov-west-1.amazonaws.com<br>athena.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

To download the latest version of the JDBC driver and its documentation, see Using Athena with the JDBC Driver.

For more information about the previous versions of the JDBC driver and their documentation, see Using the Previous Version of the JDBC Driver.

To download the latest and previous versions of the ODBC driver and their documentation, see Connecting to Athena with ODBC.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Active DDL queries | Each supported Region: 20 | Yes | The number of active DDL queries. DDL queries |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | include CREATE TABLE and ALTER TABLE ADD PARTITION queries. |
| Active DML queries | us-east-1: 200<br>us-east-2: 150<br>us-west-2: 150<br>ap-northeast-1: 150<br>ap-northeast-2: 100<br>ap-south-1: 100<br>ap-southeast-1: 100<br>ap-southeast-2: 100<br>eu-central-1: 150<br>eu-west-1: 150<br>eu-west-2: 100<br>Each of the other supported Regions: 20 | Yes | The number of active DML queries. DML queries include SELECT, CREATE TABLE AS (CTAS), and INSERT INTO queries. The specific quotas vary by AWS Region. |
| Apache Spark DPU concurrency | Each supported Region: 160 | No | The maximum number of DPUs you can consume concurrently for Apache Spark calculations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for BatchGetNamedQuery API | Each supported Region: 2 | No | Multiplier factor to scale burst capacity for BatchGetNamedQuery API based on the rate quota (Burst = Rate * BurstMultiplierQuota). BatchGetNamedQuery Burst is the maximum number of BatchGetNamedQuery API tokens you can accumulate in the account in the region |
| Burst multiplier quota for BatchGetQueryExecution API | Each supported Region: 2 | No | Multiplier factor to scale burst capacity for BatchGetQueryExecution API based on the rate quota (Burst = Rate * BurstMultiplierQuota). BatchGetQueryExecution Burst is the maximum number of BatchGetQueryExecution API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for CreateNamedQuery API | Each supported Region: 4 | No | Multiplier factor to scale burst capacity for CreateNamedQuery API based on the rate quota (Burst = Rate * BurstMultiplierQuota). CreateNamedQuery Burst is the maximum number of CreateNamedQuery API tokens you can accumulate in the account in the region |
| Burst multiplier quota for DeleteNamedQuery API | Each supported Region: 4 | No | Multiplier factor to scale burst capacity for DeleteNamedQuery API based on the rate quota (Burst = Rate * BurstMultiplierQuota). DeleteNamedQuery Burst is the maximum number of DeleteNamedQuery API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for GetNamedQuery API | Each supported Region: 4 | No | Multiplier factor to scale burst capacity for GetNamedQuery API based on the rate quota (Burst = Rate * BurstMultiplierQuota). GetNamedQuery Burst is the maximum number of GetNamedQuery API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for GetQueryExecution API | us-west-1: 2<br>af-south-1: 2<br>ap-east-1: 2<br>ap-northeast-3: 2<br>ap-south-2: 2<br>ap-southeast-3: 2<br>ap-southeast-4: 2<br>ca-central-1: 2<br>ca-west-1: 2<br>eu-central-2: 2<br>eu-north-1: 2<br>eu-south-1: 2<br>eu-south-2: 2<br>eu-west-3: 2<br>il-central-1: 2<br>sa-east-1: 2<br>Each of the other supported Regions: 1 | No | Multiplier factor to scale burst capacity for GetQueryExecution API based on the rate quota (Burst = Rate * BurstMultiplierQuota). GetQueryExecution Burst is the maximum number of GetQueryExecution API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for GetQueryResults API | us-east-1: 5<br>us-east-2: 5<br>us-west-2: 5<br>ap-northeast-1: 5<br>ap-northeast-2: 5<br>ap-south-1: 5<br>ap-southeast-1: 5<br>ap-southeast-2: 5<br>eu-central-1: 5<br>eu-west-1: 5<br>eu-west-2: 5<br>Each of the other supported Regions: 2 | No | Multiplier factor to scale burst capacity for GetQueryResults API based on the rate quota (Burst = Rate * BurstMultiplierQuota). GetQueryResults Burst is the maximum number of GetQueryResults API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for ListNamedQueries API | Each supported Region: 2 | No | Multiplier factor to scale burst capacity for ListNamedQueries API based on the rate quota (Burst = Rate * BurstMultiplierQuota). ListNamedQueries Burst is the maximum number of ListNamedQueries API tokens you can accumulate in the account in the region |
| Burst multiplier quota for ListQueryExecutions API | Each supported Region: 2 | No | Multiplier factor to scale burst capacity for ListQueryExecutions API based on the rate quota (Burst = Rate * BurstMultiplierQuota). ListQueryExecutions Burst is the maximum number of ListQueryExecutions API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst multiplier quota for StartQueryExecution API | us-east-1: 5<br>us-east-2: 5<br>us-west-2: 5<br>ap-northeast-1: 5<br>eu-central-1: 5<br>eu-west-1: 5<br>Each of the other supported Regions: 4 | No | Multiplier factor to scale burst capacity for StartQueryExecution API based on the rate quota (Burst = Rate * BurstMultiplierQuota). StartQueryExecution Burst is the maximum number of StartQueryExecution API tokens you can accumulate in the account in the region |
| Burst multiplier quota for StopQueryExecution API | Each supported Region: 4 | No | Multiplier factor to scale burst capacity for StopQueryExecution API based on the rate quota (Burst = Rate * BurstMultiplierQuota). StopQueryExecution Burst is the maximum number of StopQueryExecution API tokens you can accumulate in the account in the region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DDL query timeout | Each supported Region: 600 | No | The maximum amount of time in minutes a DDL query can run before it gets cancelled. DDL queries include CREATE TABLE and CREATE TABLE ADD PARTITION queries. |
| DML query timeout | Each supported Region: 30 | Yes | The maximum amount of time in minutes a DML query can run before it gets cancelled. DML queries include SELECT, CTAS and INSERT queries. |
| Maximum number of workgroups per account | Each supported Region: 1,000 | No | The maximum number of workgroups that can be created in an account. |
| On-demand DPUs per account | Each supported Region: 4 | Yes | The maximum number of DPUs (Data Processing Units) that can be active across all Spark workloads in this account. This limit applies only to Athena Spark Version 3.5 and does not affect Athena Engine Version 3. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Provisioned DPU | Each supported Region: 1,000 | No | Maximum number of DPUs provisioned for the account. |
| Replenishment rate of BatchGetNamedQuery API calls | Each supported Region: 5 per second | Yes | Maximum call rate per second for BatchGetNamedQuery API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of BatchGetQueryExecution API calls | Each supported Region: 10 per second | Yes | Maximum call rate per second for BatchGetQueryExecution API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |
| Replenishment rate of CreateNamedQuery API calls | Each supported Region: 5 per second | Yes | Maximum call rate per second for CreateNamedQuery API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of DeleteNamedQuery API calls | Each supported Region: 5 per second | Yes | Maximum call rate per second for DeleteNamedQuery API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |
| Replenishment rate of GetNamedQuery API calls | Each supported Region: 5 per second | Yes | Maximum call rate per second for GetNamedQuery API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of GetQueryExecution API calls | us-east-1: 500 per second<br><br>us-east-2: 500 per second<br><br>us-west-2: 500 per second<br><br>ap-northeast-1: 500 per second<br><br>ap-northeast-2: 500 per second<br><br>ap-south-1: 500 per second<br><br>ap-southeast-1: 500 per second<br><br>ap-southeast-2: 500 per second<br><br>eu-central-1: 500 per second<br><br>eu-west-1: 500 per second<br><br>eu-west-2: 500 per second<br><br>Each of the other supported Regions: 100 per second | Yes | Maximum call rate per second for GetQueryExecution API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of GetQueryResults API calls | Each supported Region: 100 per second | Yes | Maximum call rate per second for GetQueryResults API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |
| Replenishment rate of ListNamedQueries API calls | Each supported Region: 5 per second | Yes | Maximum call rate per second for ListNamedQueries API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of ListQueryExecutions API calls | Each supported Region: 5 per second | Yes | Maximum call rate per second for ListQueryExecutions API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of StartQueryExecution API calls | us-east-1: 100 per second<br><br>us-east-2: 100 per second<br><br>us-west-2: 100 per second<br><br>ap-northeast-1: 100 per second<br><br>ap-northeast-2: 100 per second<br><br>ap-south-1: 100 per second<br><br>ap-southeast-1: 100 per second<br><br>ap-southeast-2: 100 per second<br><br>eu-central-1: 100 per second<br><br>eu-west-1: 100 per second<br><br>eu-west-2: 100 per second<br><br>Each of the other supported Regions: 20 per second | Yes | Maximum call rate per second for StartQueryExecution API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of StopQuery Execution API calls | us-east-1: 50 per second<br><br>us-east-2: 50 per second<br><br>us-west-2: 50 per second<br><br>ap-northeast-1: 50 per second<br><br>ap-northeast-2: 50 per second<br><br>ap-south-1: 50 per second<br><br>ap-southeast-1: 50 per second<br><br>ap-southeast-2: 50 per second<br><br>eu-central-1: 50 per second<br><br>eu-west-1: 50 per second<br><br>eu-west-2: 50 per second<br><br>Each of the other supported Regions: 20 per second | Yes | Maximum call rate per second for StopQuery Execution API in this account in the current Region. For this API, the service adds tokens to your account at a fixed rate. One token allows you to make one API call. You accumulate tokens for an API up to a maximum of (Replenishment rate quota of the API * Burst multiplier quota of the API) |

For more information, see Service quotas in the *Amazon Athena User Guide*.

# AWS Audit Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | auditmanager.us-east-2.amazonaws.com | HTTPS |
| | | auditmanager-fips.us-east-2.api.aws | HTTPS |
| | | auditmanager-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | auditmanager.us-east-1.amazonaws.com | HTTPS |
| | | auditmanager-fips.us-east-1.api.aws | HTTPS |
| | | auditmanager-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | auditmanager.us-west-1.amazonaws.com | HTTPS |
| | | auditmanager-fips.us-west-1.api.aws | HTTPS |
| | | auditmanager-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | auditmanager.us-west-2.amazonaws.com | HTTPS |
| | | auditmanager-fips.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | auditmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | auditmanager.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | auditmanager.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | auditmanager.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | auditmanager.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | auditmanager.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | auditmanager.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | auditmanager.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | auditmanager.eu-west-2.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Accounts in scope across all assessments | Each supported Region: 250 | Yes | The maximum number of accounts in scope across all assessments per region |
| Controls per framework | Each supported Region: 1,400 | Yes | The maximum number of controls per framework per account per region |
| Custom controls | Each supported Region: 500 | Yes | The maximum number of custom controls per account per region |
| Custom frameworks | Each supported Region: 100 | Yes | The maximum number of custom frameworks per account per region |
| Running assessments | Each supported Region: 100 | Yes | The maximum number of running assessments per account per region |

# Amazon Augmented AI endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | a2i-runtime.sagemaker.us-east-2.amazonaws.com | HTTP and HTTPS | |
| US East (N. Virginia) | us-east-1 | a2i-runtime.sagemaker.us-east-1.amazonaws.com | HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | a2i-runtime.sagemaker.us-west-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | a2i-runtime.sagemaker.ap-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | a2i-runtime.sagemaker.ap-northeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | a2i-runtime.sagemaker.ap-southeast-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | a2i-runtime.sagemaker.ap-southeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | a2i-runtime.sagemaker.ap-northeast-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | a2i-runtime.sagemaker.ca-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | a2i-runtime.sagemaker.eu-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | a2i-runtime.sagemaker.eu-west-1.amazonaws.com | HTTP and HTTPS | |
| Europe (London) | eu-west-2 | a2i-runtime.sagemaker.eu-west-2.amazonaws.com | HTTP and HTTPS | |

## Service quotas

| Name | Default | Adjustable |
|---|---|---|
| Flow definitions | 100 | Yes |
| Worker task templates (human task UIs) | 100 | Yes |
| *In-flight* human loops per flow definition (private or vendor work team) † | 5,000 | No |
| *In-flight* human loops per flow definition (Amazon Mechanical Turk work team) † | 1,000 | No |

† Human loops are considered *in-flight* when their status is `InProgress` or `Stopping`.

# Amazon Aurora endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Amazon Aurora MySQL-Compatible Edition

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | rds.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | rds.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | rds.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | rds.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | rds.af-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | rds.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | rds.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rds.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rds.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rds.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rds.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | rds.ap-southeast-6.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | rds.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | rds.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | rds.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | rds.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | rds.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | rds.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | rds.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | rds.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | rds.ca-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | rds.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | rds.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | rds.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | rds.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | rds.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | rds.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | rds.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | rds.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | rds.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | rds.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | rds.me-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | rds.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | rds.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | rds.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | rds.us-gov-west-1.amazonaws.com | HTTPS | |

## Amazon Aurora PostgreSQL-Compatible Edition

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | rds.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | rds.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | rds.us-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | rds.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | rds.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | rds.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | rds.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rds.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rds.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rds.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rds.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | rds.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | rds.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | rds.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | rds.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | rds.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | rds.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | rds.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | rds.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | rds.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | rds.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | rds.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | rds.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | rds.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | rds.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | rds.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | rds.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | rds.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | rds.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | rds.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | rds.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | rds.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | rds.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | rds.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | rds.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | rds.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Authorizations per DB security group | Each supported Region: 20 | No | Number of security group authorizations per DB security group |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Custom endpoints per DB cluster | Each supported Region: 5 | Yes | The maximum number of custom endpoints that you can create per Aurora DB cluster in this account in the current Region. This value reflects the highest number of custom endpoints in a DB cluster in the account. Other DB clusters in the account might have a lower number of custom endpoints. |
| Custom engine versions | Each supported Region: 40 | Yes | The maximum number of custom engine versions allowed in this account in the current Region |
| DB cluster parameter groups | Each supported Region: 50 | Yes | The maximum number of DB cluster parameter groups |
| DB clusters | Each supported Region: 40 | Yes | The maximum number of Aurora clusters allowed in this account in the current Region |
| DB instances | ap-south-1: 20<br><br>Each of the other supported Regions: 40 | Yes | The maximum number of DB instances allowed in this account in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DB shard groups | Each supported Region: 5 | Yes | The maximum number of DB shard groups for Aurora Limitless Database in this account in the current Region |
| DB subnet groups | ap-south-1: 20<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of DB subnet groups |
| Data API HTTP request body size | Each supported Region: 4 Megabytes | No | The maximum size allowed for the HTTP request body. |
| Data API maximum concurrent cluster-secret pairs | Each supported Region: 30 | No | The maximum number of unique pairs of Aurora Serverless v1 DB clusters and secrets in concurrent Data API requests for this account in the current AWS Region. |
| Data API maximum concurrent requests | Each supported Region: 500 | No | The maximum number of Data API requests to an Aurora Serverless v1 DB cluster that use the same secret and can be processed at the same time. Additional requests are queued and processed as in-process requests complete. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Data API maximum result set size | Each supported Region: 1 Megabytes | No | The maximum size of the database result set that can be returned by the Data API. |
| Data API maximum size of JSON response string | Each supported Region: 10 Megabytes | No | The maximum size of the simplified JSON response string returned by the RDS Data API. |
| Data API requests per second | Each supported Region: 1,000 per second | No | The maximum number of requests to the Data API per second allowed for this account in the current AWS Region. This quota only applies to Amazon Aurora Serverless v1 clusters. |
| Event subscriptions | Each supported Region: 20 | Yes | The maximum number of event subscriptions |
| IAM roles per DB cluster | Each supported Region: 5 | Yes | The maximum number of IAM roles associated with a DB cluster |
| IAM roles per DB instance | Each supported Region: 5 | Yes | The maximum number of IAM roles associated with a DB instance |
| Integrations | Each supported Region: 100 | No | The maximum number of integrations allowed in this account in the current AWS Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Manual DB cluster snapshots | Each supported Region: 100 | Yes | The maximum number of manual DB cluster snapshots |
| Manual DB instance snapshots | Each supported Region: 100 | Yes | The maximum number of manual DB instance snapshots |
| Option groups | Each supported Region: 20 | Yes | The maximum number of option groups |
| Parameter groups | ap-south-1: 20<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of parameter groups |
| Proxies | Each supported Region: 20 | Yes | The maximum number of proxies allowed in this account in the current AWS Region |
| Read replicas per primary | Each supported Region: 15 | Yes | The maximum number of read replicas per primary DB instance. This quota cant be adjusted for Amazon Aurora. |
| Reserved DB instances | ap-south-1: 20<br><br>Each of the other supported Regions: 40 | Yes | The maximum number of reserved DB instances allowed in this account in the current AWS Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Security groups | ap-south-1: 20<br><br>Each of the other supported Regions: 25 | Yes | The maximum number of DB security groups |
| Subnets per DB subnet group | Each supported Region: 20 | No | The maximum number of subnets per DB subnet group |
| Total storage for all DB instances | Each supported Region: 100,000 | Yes | The maximum total storage (in GB) on EBS volumes for all Amazon RDS DB instances added together. This quota does not apply to Amazon Aurora, which has a maximum cluster volume of 128 TiB for each DB cluster. |

# Amazon Aurora DSQL endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | dsql.us-east-2.api.aws<br>dsql-fips.us-east-2.api.aws | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | dsql.us-east-1.api.aws<br>dsql-fips.us-east-1.api.aws | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | dsql.us-west-2.api.aws<br>dsql-fips.us-west-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | dsql.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | dsql.ap-southeast-4.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | dsql.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | dsql.ap-northeast-3.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | dsql.ap-northeast-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | dsql.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | dsql.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | dsql.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | dsql.ca-central-1.api.aws<br>dsql-fips.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | dsql.ca-west-1.api.aws<br>dsql-fips.ca-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | dsql.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | dsql.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | dsql.eu-west-2.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | dsql.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | dsql.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | dsql.eu-north-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | dsql.sa-east-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Cluster size | Each supported Region: 10,000 Gigabytes | Yes | The maximum capacity (in GB) for an Aurora DSQL cluster. This quota applies to both single-Region and multi-Region clusters. |
| Connections per cluster | Each supported Region: 10,000 | Yes | The maximum number of active connections per cluster that you can have in this account in the current Region. |
| Multi-Region clusters | Each supported Region: 5 | Yes | The maximum number of multi region clusters allowed per account. |
| Single-Region clusters | Each supported Region: 20 | Yes | The maximum number of active clusters allowed |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| | | | per account in a given region. |

# AWS Auto Scaling endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Note**

This page provides service endpoints and service quotas related to scaling plans, an AWS Auto Scaling resource. If you are trying to find endpoints and quotas for Amazon EC2 Auto Scaling or Application Auto Scaling, see the following:

- Amazon EC2 Auto Scaling
- Application Auto Scaling

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | autoscaling-plans.us-east-2.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | autoscaling-plans.us-east-1.amazonaws.com | HTTP and HTTPS | |
| US West (N. California) | us-west-1 | autoscaling-plans.us-west-1.amazonaws.com | HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | autoscaling-plans.us-west-2.amazonaws.com | HTTP and HTTPS | |
| Africa (Cape Town) | af-south-1 | autoscaling-plans.af-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | autoscaling-plans.ap-east-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | autoscaling-plans.ap-southeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | autoscaling-plans.ap-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | autoscaling-plans.ap-northeast-3.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | autoscaling-plans.ap-northeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | autoscaling-plans.ap-southeast-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | autoscaling-plans.ap-southeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | autoscaling-plans.ap-northeast-1.amazonaws.com | HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | autoscaling-plans.ca-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | autoscaling-plans.eu-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | autoscaling-plans.eu-west-1.amazonaws.com | HTTP and HTTPS | |
| Europe (London) | eu-west-2 | autoscaling-plans.eu-west-2.amazonaws.com | HTTP and HTTPS | |
| Europe (Milan) | eu-south-1 | autoscaling-plans.eu-south-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | autoscaling-plans.eu-west-3.amazonaws.com | HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | autoscaling-plans.eu-north-1.amazonaws.com | HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | autoscaling-plans.me-south-1.amazonaws.com | HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | autoscaling-plans.sa-east-1.amazonaws.com | HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | autoscaling-plans.us-gov-east-1.amazonaws.com | HTTP and HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | autoscaling-plans.us-gov-west-1.amazonaws.com | HTTP and HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Scaling instructions per scaling plan | Each supported Region: 500 | No | The maximum number of scaling instructions per scaling plan. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Scaling plans | Each supported Region: 100 | Yes | The maximum number of scaling plans that you can create in this account in the current Region. A scaling plan tells AWS Auto Scaling how to scale a collection of resources. |
| Target tracking configurations per scaling instruction | Each supported Region: 10 | No | The maximum number of target tracking configurations per scaling instruction. |

For more information, see Quotas for your scaling plans in the *AWS Auto Scaling User Guide*.

# AWS MCP Server endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The following table shows the AWS Regions and endpoints for AWS MCP Server.

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | aws-mcp.us-east-1.api.aws/mcp | HTTPS |
| Europe (Frankfurt) | eu-central-1 | aws-mcp.eu-central-1.api.aws/mcp | HTTPS |

### Service quotas

For more information, see Quotas for AWS MCP Server in the *Agent Toolkit User Guide*.

## AWS B2B Data Interchange endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

### Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | b2bi.us-east-2.amazonaws.com | HTTPS |
| | | b2bi-fips.us-east-2.api.aws | HTTPS |
| | | b2bi-fips.us-east-2.amazonaws.com | HTTPS |
| | | b2bi.us-east-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | b2bi.us-east-1.amazonaws.com | HTTPS |
| | | b2bi-fips.us-east-1.api.aws | HTTPS |
| | | b2bi-fips.us-east-1.amazonaws.com | HTTPS |
| | | b2bi.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | b2bi.us-west-2.amazonaws.com | HTTPS |
| | | b2bi-fips.us-west-2.api.aws | HTTPS |
| | | b2bi-fips.us-west-2.amazonaws.com | HTTPS |
| | | b2bi.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | b2bi.ap-south-2.amazonaws.com | HTTPS |
| | | b2bi.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | b2bi.ap-southeast-2.amazonaws.com | HTTPS |
| | | b2bi.ap-southeast-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | b2bi.ca-central-1.amazonaws.com | HTTPS |
| | | b2bi-fips.ca-central-1.api.aws | HTTPS |
| | | b2bi-fips.ca-central-1.amazonaws.com | HTTPS |
| | | b2bi.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | b2bi.eu-central-1.amazonaws.com | HTTPS |
| | | b2bi.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | b2bi.eu-west-1.amazonaws.com<br><br>b2bi.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Paris) | eu-west-3 | b2bi.eu-west-3.amazonaws.com<br><br>b2bi.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |

### Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Capabilities | Each supported Region: 100 | Yes | Maximum number of Capabilities. |
| Partnerships | Each supported Region: 700 | Yes | Maximum number of Partnerships. |
| Profiles | Each supported Region: 5 | Yes | Maximum number of Profiles. |
| Transformers | Each supported Region: 500 | Yes | Maximum number of Transformers. |

## AWS Backup endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | backup.us-east-2.amazonaws.com<br>backup-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | backup.us-east-1.amazonaws.com<br>backup-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | backup.us-west-1.amazonaws.com<br>backup-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | backup.us-west-2.amazonaws.com<br>backup-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | backup.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | backup.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | backup.ap-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | backup.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | backup.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | backup.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | backup.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | backup.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | backup.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | backup.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | backup.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | backup.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | backup.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | backup.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | backup.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | backup.ca-central-1.amazonaws.com<br>backup-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | backup.ca-west-1.amazonaws.com<br>backup-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | backup.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | backup.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | backup.eu-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | backup.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | backup.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | backup.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | backup.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | backup.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | backup.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | backup.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | backup.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | backup.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | backup.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | backup.us-gov-east-1.amazonaws.com<br>backup-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | backup.us-gov-west-1.amazonaws.com<br>backup-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Backup plans per Region per account | Each supported Region: 300 | Yes | Number of backup plans in this account in the current Region |
| Backup vaults per Region per account | Each supported Region: 300 | Yes | Number of backup vaults in this account in the current Region |
| Concurrent backup copies per supported service per account | Each supported Region: 5 | No | Number of concurrent backup copies per supported service per account |
| Concurrent backup jobs per resource | Each supported Region: 1 | No | Number of concurrent backup jobs per resource |
| Framework controls per Region per account | Each supported Region: 50 | Yes | Number of framework controls in this account in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Frameworks per Region per account | Each supported Region: 10 | Yes | Number of frameworks in this account in the current Region |
| Frameworks per report plan | Each supported Region: 1,000 | No | Number of frameworks per report plan |
| Maximum backup nest level | Each supported Region: 10 | Yes | Maximum number of nested levels allowed for backups |
| Metadata tags per backup | Each supported Region: 50 | No | Number of metadata tags per backup |
| Recovery points per backup vault | Each supported Region: 1,000,000 | Yes | Number of recovery points per backup vault |
| Report plans per Region per account | Each supported Region: 20 | Yes | Number of report plans in this account in the current Region |
| Versions per backup plan | Each supported Region: 2,000 | Yes | Number of versions per backup plan |

If you regularly receive throttling exceptions, consider using a rate limiter.

| API name | Default calls/sec |
|---|---|
| CreateBackupPlan \| CreateBackupSelection \| DeleteBackupPlan \| DeleteBackupSelection \| DeleteBackupVault \| DeleteBackupVaultAccessPolicy \| DeleteBackupVaultNotifications \| DescribeBackupVault \| ExportBackupPlanTemplate \| GetBackupPlanFromJSON \| GetBackupPlanFromTemplate \| PutBackupVaultNotifications \| StartBackupJob | 5 |

| API name | Default calls/sec |
| --- | --- |
| \| StartRestoreJob \| StopBackupJob \| TagResource \| UntagResource \| UpdateBackupPlan \| UpdateRecoveryPointLifecycle | |
| DeleteRecoveryPoint \| DescribeProtectedResource | 10 |
| DescribeBackupJob \| DescribeRecoveryPoint \| DescribeRestoreJob \| GetBackupPlan \| GetBackupSelection \| GetBackupVaultAccessPolicy \| GetBackupVaultNotifications \| GetRecoveryPointRestoreMetadata \| GetSupportedResourceTypes | 15 |
| ListBackupJobs \| ListBackupPlans \| ListBackupPlanTemplates \| ListBackupPlanVersions \| ListBackupSelections \| ListBackupVaults \| ListProtectedResources \| ListRecoveryPointByResource \| ListRecoveryPointsByBackupVault \| ListRecoveryPointsByResource \| ListRestoreJobs \| ListTags | 20 |
| Sum of All API Calls | 50 |

For additional information, see Quotas in the *AWS Backup Developer Guide*.

## AWS Batch endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | batch.us-east-2.amazonaws.com | HTTPS | |
| | | fips.batch.us-east-2.amazonaws.com | HTTPS | |
| | | batch-fips.us-east-2.api.aws | HTTPS | |
| | | batch.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | batch.us-east-1.amazonaws.com | HTTPS | |
| | | fips.batch.us-east-1.amazonaws.com | HTTPS | |
| | | batch-fips.us-east-1.api.aws | HTTPS | |
| | | batch.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | batch.us-west-1.amazonaws.com | HTTPS | |
| | | fips.batch.us-west-1.amazonaws.com | HTTPS | |
| | | batch-fips.us-west-1.api.aws | HTTPS | |
| | | batch.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | batch.us-west-2.amazonaws.com | HTTPS | |
| | | fips.batch.us-west-2.amazonaws.com | HTTPS | |
| | | batch-fips.us-west-2.api.aws | HTTPS | |
| | | batch.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | batch.af-south-1.amazonaws.com | HTTPS | |
| | | batch.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | batch.ap-east-1.amazonaws.com<br>batch.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | batch.ap-south-2.amazonaws.com<br>batch.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | batch.ap-southeast-3.amazonaws.com<br>batch.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | batch.ap-southeast-5.amazonaws.com<br>batch.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | batch.ap-southeast-4.amazonaws.com<br>batch.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | batch.ap-south-1.amazonaws.com<br>batch.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | batch.ap-southeast-6.amazonaws.com<br>batch.ap-southeast-6.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | batch.ap-northeast-3.amazonaws.com<br>batch.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | batch.ap-northeast-2.amazonaws.com<br>batch.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | batch.ap-southeast-1.amazonaws.com<br>batch.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | batch.ap-southeast-2.amazonaws.com<br>batch.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | batch.ap-east-2.amazonaws.com<br>batch.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | batch.ap-southeast-7.amazonaws.com<br>batch.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | batch.ap-northeast-1.amazonaws.com<br>batch.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | batch.ca-central-1.amazonaws.com<br>batch.ca-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | batch.ca-west-1.amazonaws.com | HTTPS |
| | | batch.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | batch.eu-central-1.amazonaws.com | HTTPS |
| | | batch.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | batch.eu-west-1.amazonaws.com | HTTPS |
| | | batch.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | batch.eu-west-2.amazonaws.com | HTTPS |
| | | batch.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | batch.eu-south-1.amazonaws.com | HTTPS |
| | | batch.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | batch.eu-west-3.amazonaws.com | HTTPS |
| | | batch.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | batch.eu-south-2.amazonaws.com | HTTPS |
| | | batch.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | batch.eu-north-1.amazonaws.com | HTTPS |
| | | batch.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | batch.eu-central-2.amazonaws.com | HTTPS |
| | | batch.eu-central-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | batch.il-central-1.amazonaws.com | HTTPS |
| | | batch.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | batch.mx-central-1.amazonaws.com | HTTPS |
| | | batch.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | batch.me-south-1.amazonaws.com | HTTPS |
| | | batch.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | batch.me-central-1.amazonaws.com | HTTPS |
| | | batch.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | batch.sa-east-1.amazonaws.com | HTTPS |
| | | batch.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | batch.us-gov-east-1.amazonaws.com | HTTPS |
| | | batch-fips.us-gov-east-1.api.aws | HTTPS |
| | | batch.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | batch.us-gov-west-1.amazonaws.com | HTTPS |
| | | batch-fips.us-gov-west-1.api.aws | HTTPS |
| | | batch.us-gov-west-1.api.aws | HTTPS |

# Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Compute environment limit | Each supported Region: 50 | No | Maximum number of compute environments per account, per region. |
| Compute environments per job queue limit. | Each supported Region: 3 | No | Maximum number of compute environments per job queue. |
| Job definition size limit | Each supported Region: 24 | No | Maximum job definition size (for RegisterJobDefinition API operations), measured in KiB. |
| Job dependencies limit | Each supported Region: 20 | No | Maximum number of job dependencies per job. |
| Job payload size limit | Each supported Region: 30 | No | Maximum job payload size (for SubmitJob API operations), measured in KiB. |
| Job queue limit | Each supported Region: 50 | No | Maximum number of job queues per account, per region. |
| Maximum array size limit | Each supported Region: 10,000 | No | Maximum array size for array jobs. |
| Service Job serviceRequestPayload size | Each supported Region: 10 | No | Maximum serviceRequestpayload size (within SubmitServiceJob API |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | operations), measured in KiB. |
| Service Job total payload size | Each supported Region: 30 | No | Maximum total job payload size (for SubmitServiceJob API operations), measured in KiB. |
| Service environment | Each supported Region: 50 | No | Maximum number of service environments per account, per region. |
| Service environments per job queue | Each supported Region: 1 | No | Maximum number of service environments per job queue. |
| Share identifiers per job queue limit. | Each supported Region: 500 | No | Maximum number of share identifiers per job queue. |
| Submitted state jobs limit | Each supported Region: 1,000,000 | No | Maximum number of jobs in SUBMITTED state. |
| Transactions per second for SubmitJob limit | Each supported Region: 50 | No | Maximum number of transactions per second (TPS) for each account for SubmitJob operations. |
| Transactions per second for SubmitServiceJob | Each supported Region: 5 | No | Maximum number of transactions per second (TPS) for each account for SubmitServiceJob operations. |

For more information, see Service Quotas in the *AWS Batch User Guide*.

# Amazon Bedrock endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Amazon Bedrock service endpoints

### Amazon Bedrock control plane APIs

The following table provides a list of Region-specific endpoints that Amazon Bedrock supports for managing, training, and deploying models. Use these endpoints for Amazon Bedrock API operations.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | bedrock.us-east-2.amazonaws.com | HTTPS | |
| | | bedrock-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | bedrock.us-east-1.amazonaws.com | HTTPS | |
| | | bedrock-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | bedrock.us-west-1.amazonaws.com | HTTPS | |
| | | bedrock-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | bedrock.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | bedrock-fips.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | bedrock.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | bedrock.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | bedrock.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | bedrock.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | bedrock.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | bedrock.ap-southeast-6.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | bedrock.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | bedrock.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | bedrock.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | bedrock.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | bedrock.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | bedrock.ca-central-1.amazonaws.com<br>bedrock-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | bedrock.ca-west-1.amazonaws.com<br>bedrock-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | bedrock.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | bedrock.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | bedrock.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | bedrock.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | bedrock.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | bedrock.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | bedrock.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | bedrock.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | bedrock.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | bedrock.mx-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | bedrock.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | bedrock.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | bedrock.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | bedrock.us-gov-east-1.amazonaws.com<br>bedrock-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | bedrock.us-gov-west-1.amazonaws.com<br>bedrock-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

## Amazon Bedrock runtime APIs

The following table provides a list of Region-specific endpoints that Amazon Bedrock supports for making inference requests for models hosted in Amazon Bedrock. Use these endpoints for Amazon Bedrock Runtime API operations.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | bedrock-runtime.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | bedrock-runtime-fips.us-east-2.amazonaws.com | | |
| US East (N. Virginia) | us-east-1 | bedrock-runtime.us-east-1.amazonaws.com<br>bedrock-runtime-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | bedrock-runtime.us-west-2.amazonaws.com<br>bedrock-runtime-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | bedrock-runtime.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock-runtime.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | bedrock-runtime.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | bedrock-runtime.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | bedrock-runtime.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock-runtime.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock-runtime.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | bedrock-runtime.ca-central-1.amazonaws.com<br>bedrock-runtime-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | bedrock-runtime.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | bedrock-runtime.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | bedrock-runtime.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | bedrock-runtime.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | bedrock-runtime.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | bedrock-runtime.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | bedrock-runtime.eu-north-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | bedrock-runtime.eu-central-2.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | bedrock-runtime.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | bedrock-runtime.us-gov-east-1.amazonaws.com<br>bedrock-runtime-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | bedrock-runtime.us-gov-west-1.amazonaws.com<br>bedrock-runtime-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

## Agents for Amazon Bedrock build-time APIs

The following table provides a list of Region-specific endpoints that Agents for Amazon Bedrock supports for creating and managing agents and knowledge bases. Use these endpoints for Agents for Amazon Bedrock API operations.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | bedrock-agent.us-east-1.amazonaws.com | HTTPS |
| | | bedrock-agent-fips.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | bedrock-agent.us-west-2.amazonaws.com | HTTPS |
| | | bedrock-agent-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | bedrock-agent.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock-agent.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock-agent.ap-northeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | bedrock-agent.ap-northeast-2.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | bedrock-agent.ca-central-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | bedrock-agent.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | bedrock-agent.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | bedrock-agent.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | bedrock-agent.eu-west-3.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock-agent.ap-south-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | bedrock-agent.sa-east-1.amazonaws.com | HTTPS |

## Agents for Amazon Bedrock runtime APIs

The following table provides a list of Region-specific endpoints that Agents for Amazon Bedrock supports for invoking agents and querying knowledge bases. Use these endpoints for Agents for Amazon Bedrock Runtime API operations.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | bedrock-agent-runtime.us-east-1.amazonaws.com | HTTPS |
| | | bedrock-agent-runtime-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | bedrock-agent-runtime.us-west-2.amazonaws.com | HTTPS |
| | | bedrock-agent-runtime-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | bedrock-agent-runtime.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock-agent-runtime.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock-agent-runtime.ap-northeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | bedrock-agent-runtime.ap-northeast-2.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | bedrock-agent-runtime.ca-central-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | bedrock-agent-runtime.eu-central-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | bedrock-agent-runtime.eu-west-3.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | bedrock-agent-runtime.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | bedrock-agent-runtime.eu-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock-agent-runtime.ap-south-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | bedrock-agent-runtime.sa-east-1.amazonaws.com | HTTPS |

## Amazon Bedrock Data Automation APIs

The following table provides a list of Region-specific endpoints that Data Automation for Amazon Bedrock supports. Endpoints that use the word `runtime` invoke blueprints and projects to extract information from files. Use these endpoints for Amazon Bedrock Data Automation Runtime API operations. Endpoints without `runtime` are used to create blueprints and projects to provide extraction guidance. Use these endpoints for Amazon Bedrock Data Automation API Buildtime operations

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | bedrock-data-automation.us-east-2.amazonaws.com | HTTPS |
| | | bedrock-data-automation-runtime.us-east-2.amazonaws.com | HTTPS |
| | | bedrock-data-automation-fips.us-east-2.amazonaws.com | HTTPS |
| | | bedrock-data-automation-runtime-fips.us-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | bedrock-data-automation.us-east-1.amazonaws.com<br>bedrock-data-automation-runtime.us-east-1.api.aws<br>bedrock-data-automation-runtime.us-east-1.amazonaws.com<br>bedrock-data-automation.us-east-1.api.aws<br>bedrock-data-automation-fips.us-east-1.amazonaws.com<br>bedrock-data-automation-runtime-fips.us-east-1.api.aws<br>bedrock-data-automation-runtime-fips.us-east-1.amazonaws.com<br>bedrock-data-automation-fips.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | bedrock-data-automation.us-west-2.amazonaws.com<br>bedrock-data-automation-runtime.us-west-2.api.aws<br>bedrock-data-automation-runtime.us-west-2.amazonaws.com<br>bedrock-data-automation.us-west-2.api.aws<br>bedrock-data-automation-fips.us-west-2.amazonaws.com<br>bedrock-data-automation-runtime-fips.us-west-2.api.aws<br>bedrock-data-automation-runtime-fips.us-west-2.amazonaws.com<br>bedrock-data-automation-fips.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock-data-automation.ap-south-1.amazonaws.com<br>bedrock-data-automation-runtime.ap-south-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock-data-automation.ap-southeast-2.amazonaws.com<br>bedrock-data-automation-runtime.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock-data-automation.ap-northeast-1.amazonaws.com<br><br>bedrock-data-automation-runtime.ap-northeast-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | bedrock-data-automation.ca-central-1.amazonaws.com<br><br>bedrock-data-automation-runtime.ca-central-1.amazonaws.com<br><br>bedrock-data-automation-fips.ca-central-1.amazonaws.com<br><br>bedrock-data-automation-runtime-fips.ca-central-1.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | bedrock-data-automation.eu-central-1.amazonaws.com<br><br>bedrock-data-automation-runtime.eu-central-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | bedrock-data-automation.eu-west-1.amazonaws.com<br><br>bedrock-data-automation-runtime.eu-west-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| Europe (London) | eu-west-2 | bedrock-data-automation.eu-west-2.amazonaws.com<br><br>bedrock-data-automation-runtime.eu-west-2.amazonaws.com | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | bedrock-data-automation.eu-south-2.amazonaws.com<br><br>bedrock-data-automation-runtime.eu-south-2.amazonaws.com | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | bedrock-data-automation.us-gov-west-1.amazonaws.com<br><br>bedrock-data-automation-runtime.us-gov-west-1.amazonaws.com<br><br>bedrock-data-automation-fips.us-gov-west-1.amazonaws.com<br><br>bedrock-data-automation-runtime-fips.us-gov-west-1.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS<br><br>HTTPS | |

## Amazon Bedrock service quotas

**Tip**

Because Amazon Bedrock has a large number of quotas, we recommend that you view the service quotas using the console instead of using the table below. Open Amazon Bedrock quotas. Use the search field to find specific quotas. Use the gear icon to control how many quotas are displayed on a page.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Advanced Prompt Optimization) Active jobs per account | Each supported Region: 20 | Yes | The maximum number of active Advanced Prompt |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | Optimization (APO) jobs per account. |
| (Advanced Prompt Optimization) Inactive jobs per account | Each supported Region: 5,000 | Yes | The maximum number of inactive Advanced Prompt Optimization (APO) jobs per account. |
| (Automated Reasoning) Annotations in policy | Each supported Region: 10 | No | The maximum number of annotations in an Automated Reasoning policy. |
| (Automated Reasoning) CancelAutomatedReasoningPolicyBuildWorkflow requests per second | Each supported Region: 5 | Yes | The maximum number of CancelAutomatedReasoningPolicyBuildWorkflow API requests per second. |
| (Automated Reasoning) Concurrent builds per policy | Each supported Region: 2 | No | The maximum number of concurrent builds per Automated Reasoning policy. |
| (Automated Reasoning) Concurrent policy builds per account | Each supported Region: 5 | No | The maximum number of concurrent Automated Reasoning policy builds in one account. |
| (Automated Reasoning) CreateAutomatedReasoningPolicy requests per second | Each supported Region: 5 | Yes | The maximum number of CreateAutomatedReasoningPolicy API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Automated Reasoning) CreateAutomatedReasoningPolicyTestCase requests per second | Each supported Region: 5 | Yes | The maximum number of CreateAutomatedReasoningPolicyTestCase API requests per second. |
| (Automated Reasoning) CreateAutomatedReasoningPolicyVersion requests per second | Each supported Region: 5 | Yes | The maximum number of CreateAutomatedReasoningPolicyVersion API requests per second. |
| (Automated Reasoning) DeleteAutomatedReasoningPolicy requests per second | Each supported Region: 5 | Yes | The maximum number of DeleteAutomatedReasoningPolicy API requests per second. |
| (Automated Reasoning) DeleteAutomatedReasoningPolicyBuildWorkflow requests per second | Each supported Region: 5 | Yes | The maximum number of DeleteAutomatedReasoningPolicyBuildWorkflow API requests per second. |
| (Automated Reasoning) DeleteAutomatedReasoningPolicyTestCase requests per second | Each supported Region: 5 | Yes | The maximum number of DeleteAutomatedReasoningPolicyTestCase API requests per second. |
| (Automated Reasoning) ExportAutomatedReasoningPolicyVersion requests per second | Each supported Region: 5 | Yes | The maximum number of ExportAutomatedReasoningPolicyVersion API requests per second. |
| (Automated Reasoning) GetAutomatedReasoningPolicy requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicy API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Automated Reasoning) GetAutomatedReasoningPolicyAnnotations requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicyAnnotations API requests per second. |
| (Automated Reasoning) GetAutomatedReasoningPolicyBuildWorkflow requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicyBuildWorkflow API requests per second. |
| (Automated Reasoning) GetAutomatedReasoningPolicyBuildWorkflowResultAssets requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicyBuildWorkflowResultAssets API requests per second. |
| (Automated Reasoning) GetAutomatedReasoningPolicyNextScenario requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicyNextScenario API requests per second. |
| (Automated Reasoning) GetAutomatedReasoningPolicyTestCase requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicyTestCase API requests per second. |
| (Automated Reasoning) GetAutomatedReasoningPolicyTestResult requests per second | Each supported Region: 10 | Yes | The maximum number of GetAutomatedReasoningPolicyTestResult API requests per second. |
| (Automated Reasoning) ListAutomatedReasoningPolicies requests per second | Each supported Region: 5 | Yes | The maximum number of ListAutomatedReasoningPolicies API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Automated Reasoning) ListAutomatedReasoningPolicyBuildWorkflows requests per second | Each supported Region: 5 | Yes | The maximum number of ListAutomatedReasoningPolicyBuildWorkflows API requests per second. |
| (Automated Reasoning) ListAutomatedReasoningPolicyTestCases requests per second | Each supported Region: 5 | Yes | The maximum number of ListAutomatedReasoningPolicyTestCases API requests per second. |
| (Automated Reasoning) ListAutomatedReasoningPolicyTestResults requests per second | Each supported Region: 5 | Yes | The maximum number of ListAutomatedReasoningPolicyTestResults API requests per second. |
| (Automated Reasoning) Policies per account | Each supported Region: 100 | No | The maximum number of Automated Reasoning policies in one account. |
| (Automated Reasoning) Rules in policy | Each supported Region: 500 | No | The maximum number of rules in an Automated Reasoning policy. |
| (Automated Reasoning) Source document size (MB) | Each supported Region: 5 | No | The maximum source document size (MB) for creating an Automated Reasoning policy. |
| (Automated Reasoning) Source document tokens | Each supported Region: 122,880 | No | The maximum number of tokens allowed in a source document when creating an Automated Reasoning policy. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Automated Reasoning) StartAutomatedReasoningPolicyBuildWorkflow requests per second | Each supported Region: 1 | Yes | The maximum number of StartAutomatedReasoningPolicyBuildWorkflow API requests per second. |
| (Automated Reasoning) StartAutomatedReasoningPolicyTestWorkflow requests per second | Each supported Region: 1 | Yes | The maximum number of StartAutomatedReasoningPolicyTestWorkflow API requests per second. |
| (Automated Reasoning) Tests per policy | Each supported Region: 100 | No | The maximum number of tests per Automated Reasoning policy. |
| (Automated Reasoning) Types per policy | Each supported Region: 50 | No | The maximum number of types in an Automated Reasoning policy. |
| (Automated Reasoning) UpdateAutomatedReasoningPolicy requests per second | Each supported Region: 5 | Yes | The maximum number of UpdateAutomatedReasoningPolicy API requests per second. |
| (Automated Reasoning) UpdateAutomatedReasoningPolicyAnnotations requests per second | Each supported Region: 5 | Yes | The maximum number of UpdateAutomatedReasoningPolicyAnnotations API requests per second. |
| (Automated Reasoning) UpdateAutomatedReasoningPolicyTestCase requests per second | Each supported Region: 5 | Yes | The maximum number of UpdateAutomatedReasoningPolicyTestCase API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Automated Reasoning) Values per type in policy | Each supported Region: 50 | No | The maximum number of values per type in an Automated Reasoning policy. |
| (Automated Reasoning) Variables in policy | Each supported Region: 200 | No | The maximum number of variables in an Automated Reasoning policy. |
| (Automated Reasoning) Versions per policy | Each supported Region: 1,000 | No | The maximum number of versions per Automated Reasoning policy. |
| (Data Automation) (Console) Maximum document file size (MB) | Each supported Region: 200 | No | The maximum file size for console |
| (Data Automation) (Console) Maximum number of pages per document file | Each supported Region: 20 | No | The maximum number of pages per document in console |
| (Data Automation) CreateBlueprint - Max number of blueprints per account | Each supported Region: 350 | Yes | The maximum number of blueprints per account |
| (Data Automation) CreateBlueprintVersion - Max number of Blueprint versions per Blueprint | Each supported Region: 10 | Yes | The maximum number of versions per blueprint |
| (Data Automation) CreateDataAutomationLibrary - Max number of data automation libraries per account | Each supported Region: 10 | Yes | The maximum number of data automation libraries per account |
| (Data Automation) Description length for fields (Characters) | Each supported Region: 300 | No | The maximum length of description for fields in characters |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Data Automation) InvokeBlueprintOptimizationAsync - Max number of blueprint optimization concurrent jobs | Each supported Region: 3 | Yes | The maximum number of Invoke Blueprint Optimization Async open jobs |
| (Data Automation) InvokeBlueprintOptimizationAsync - Max number of blueprint optimization jobs per day | Each supported Region: 30 | No | The maximum number of Invoke Blueprint Optimization Async jobs per day |
| (Data Automation) InvokeDataAutomation(Sync) - Document - Max number of requests | Each supported Region: 60 | Yes | The maximum number of InvokeDataAutomation requests per minute for document modality |
| (Data Automation) InvokeDataAutomationAsync - Audio - Max number of concurrent jobs | us-east-1: 20<br>us-west-2: 20<br>Each of the other supported Regions: 2 | Yes | The maximum number of Invoke Data Automation Async open jobs for audios |
| (Data Automation) InvokeDataAutomationAsync - Document - Max number of concurrent jobs | ap-south-1: 5<br>ca-central-1: 5<br>eu-south-2: 5<br>eu-west-2: 5<br>Each of the other supported Regions: 25 | Yes | The maximum number of Invoke Data Automation Async open jobs for documents |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Data Automation) InvokeDataAutomationAsync - Image - Max number of concurrent jobs | us-east-1: 20<br>us-west-2: 20<br>Each of the other supported Regions: 5 | Yes | The maximum number of Invoke Data Automation Async open jobs for images |
| (Data Automation) InvokeDataAutomationAsync - Max number of open jobs | Each supported Region: 1,800 | No | The maximum number of Invoke Data Automation Async open jobs for images |
| (Data Automation) InvokeDataAutomationAsync - Video - Max number of concurrent jobs | us-east-1: 20<br>us-west-2: 20<br>Each of the other supported Regions: 3 | Yes | The maximum number of Invoke Data Automation Async open jobs for videos |
| (Data Automation) Max number of vocabulary phrases per library | Each supported Region: 500 | Yes | The maximum number of custom vocabulary phrases that can be configured per library |
| (Data Automation) Maximum Audio Sample Rate (Hz) | Each supported Region: 48,000 | No | The maximum audio sample rate |
| (Data Automation) Maximum Blueprints per Project (Audios) | Each supported Region: 1 | No | The maximum number of blueprints per projects for audios |
| (Data Automation) Maximum Blueprints per Project (Documents) | Each supported Region: 40 | No | The maximum number of blueprints per projects for documents |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| (Data Automation) Maximum Blueprints per Project (Images) | Each supported Region: 1 | No | The maximum number of blueprints per projects for images |
| (Data Automation) Maximum Blueprints per Project (Videos) | Each supported Region: 1 | No | The maximum number of blueprints per projects for videos |
| (Data Automation) Maximum JSON Blueprint Size (Characters) | Each supported Region: 100,000 | No | The maximum size of JSON in Characters |
| (Data Automation) Maximum Levels of Field Hierarchy | Each supported Region: 1 | No | The maximum number level of field hierarchy |
| (Data Automation) Maximum Number of pages per document | Each supported Region: 3,000 | No | The maximum number of pages per document |
| (Data Automation) Maximum Resolution | Each supported Region: 8,000 | No | The maximum resolution for Images |
| (Data Automation) Maximum audio file size (MB) | Each supported Region: 2,048 | No | The maximum file size for Audio |
| (Data Automation) Maximum audio length (Minutes) | Each supported Region: 240 | No | The maximum length for audio in minutes |
| (Data Automation) Maximum document file size (MB) | Each supported Region: 500 | No | The maximum file size |
| (Data Automation) Maximum image file size (MB) | Each supported Region: 5 | No | The maximum file size for Images |
| (Data Automation) Maximum instruction field length for Audio Blueprint - (Characters) | Each supported Region: 500 | Yes | The maximum length of instruction field for audio blueprint in characters |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Data Automation) Maximum number of Blueprints per Start Inference request (Audios) | Each supported Region: 1 | No | The maximum number of inline blueprint in Start inference request |
| (Data Automation) Maximum number of Blueprints per Start Inference request (Documents) | Each supported Region: 10 | No | The maximum number of inline blueprint in Start inference request |
| (Data Automation) Maximum number of Blueprints per Start Inference request (Images) | Each supported Region: 1 | No | The maximum number of inline blueprint in Start inference request |
| (Data Automation) Maximum number of Blueprints per Start Inference request (Videos) | Each supported Region: 1 | No | The maximum number of inline blueprint in Start inference request |
| (Data Automation) Maximum number of list fields per Blueprint | Each supported Region: 15 | No | The maximum number of list fields per Blueprint |
| (Data Automation) Maximum video file size (MB) | Each supported Region: 10,240 | No | The maximum file size for Videos |
| (Data Automation) Maximum video length (Minutes) | Each supported Region: 240 | No | The maximum length for videos in minutes |
| (Data Automation) Minimum Audio Sample Rate (Hz) | Each supported Region: 8,000 | No | The minimum audio sample rate |
| (Data Automation) Minimum audio length (Miliseconds) | Each supported Region: 500 | No | The minimum length for audio in miliseconds |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Evaluation) Number of concurrent automatic model evaluation jobs | Each supported Region: 20 | No | The maximum number of automatic model evaluation jobs that you can specify at one time in this account in the current Region. |
| (Evaluation) Number of concurrent model evaluation jobs that use human workers | Each supported Region: 10 | No | The maximum number of model evaluation jobs that use human workers you can specify at one time in this account in the current Region. |
| (Evaluation) Number of custom metrics | Each supported Region: 10 | No | The maximum number of custom metrics that you can specify in a model evaluation job that uses human workers. |
| (Evaluation) Number of custom prompt datasets in a human-based model evaluation job | Each supported Region: 1 | No | The maximum number of custom prompt datasets that you can specify in a human-based model evaluation job in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Evaluation) Number of datasets per job | Each supported Region: 5 | No | The maximum number of datasets that you can specify in an automated model evaluation job. This includes both custom and built-in prompt datasets. |
| (Evaluation) Number of evaluation jobs | Each supported Region: 5,000 | No | The maximum number of model evaluation jobs that you can create in this account in the current Region. |
| (Evaluation) Number of metrics per dataset | Each supported Region: 3 | No | The maximum number of metrics that you can specify per dataset in an automated model evaluation job. This includes both custom and built-in metrics. |
| (Evaluation) Number of models in a model evaluation job that uses human workers | Each supported Region: 2 | No | The maximum number of models that you can specify in a model evaluation job that uses human workers. |
| (Evaluation) Number of models in automated model evaluation job | Each supported Region: 1 | No | The maximum number of models that you can specify in an automated model evaluation job. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Evaluation) Number of prompts in a custom prompt dataset | Each supported Region: 1,000 | No | The maximum number of prompts a custom prompt dataset can contain. |
| (Evaluation) Size of prompt | Each supported Region: 4 | No | The maximum size (in KB) of an individual prompt in a custom prompt dataset. |
| (Evaluation) Task time for workers | Each supported Region: 30 | No | The maximum length (in days) of time that a worker can have to complete tasks. |
| (Flows) Agent nodes per flow | Each supported Region: 20 | No | The maximum number of agent nodes. |
| (Flows) Collector nodes per flow | Each supported Region: 1 | No | The maximum number of collector nodes. |
| (Flows) Condition nodes per flow | Each supported Region: 5 | No | The maximum number of condition nodes. |
| (Flows) Conditions per condition node | Each supported Region: 5 | No | The maximum number of conditions per condition node. |
| (Flows) CreateFlow requests per second | Each supported Region: 2 | No | The maximum number of CreateFlow requests per second. |
| (Flows) CreateFlowAlias requests per second | Each supported Region: 2 | No | The maximum number of CreateFlowAlias requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Flows) CreateFlowVersion requests per second | Each supported Region: 2 | No | The maximum number of CreateFlowVersion requests per second. |
| (Flows) DeleteFlow requests per second | Each supported Region: 2 | No | The maximum number of DeleteFlow requests per second. |
| (Flows) DeleteFlowAlias requests per second | Each supported Region: 2 | No | The maximum number of DeleteFlowAlias requests per second. |
| (Flows) DeleteFlowVersion requests per second | Each supported Region: 2 | No | The maximum number of DeleteFlowVersion requests per second. |
| (Flows) Flow aliases per flow | Each supported Region: 10 | No | The maximum number of flow aliases. |
| (Flows) Flow executions per account | Each supported Region: 1,000 | Yes | The maximum number of flow executions per account. |
| (Flows) Flow versions per flow | Each supported Region: 10 | No | The maximum number of flow versions. |
| (Flows) Flows per account | Each supported Region: 100 | Yes | The maximum number of flows per account. |
| (Flows) GetFlow requests per second | Each supported Region: 10 | No | The maximum number of GetFlow requests per second. |
| (Flows) GetFlowAlias requests per second | Each supported Region: 10 | No | The maximum number of GetFlowAlias requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Flows) GetFlowVersion requests per second | Each supported Region: 10 | No | The maximum number of GetFlowVersion requests per second. |
| (Flows) Inline code nodes per flow | Each supported Region: 5 | No | The maximum number of inline code nodes per flow. |
| (Flows) Input nodes per flow | Each supported Region: 1 | No | The maximum number of flow input nodes. |
| (Flows) Iterator nodes per flow | Each supported Region: 1 | No | The maximum number of iterator nodes. |
| (Flows) Knowledge base nodes per flow | Each supported Region: 20 | No | The maximum number of knowledge base nodes. |
| (Flows) Lambda function nodes per flow | Each supported Region: 20 | No | The maximum number of Lambda function nodes. |
| (Flows) Lex nodes per flow | Each supported Region: 5 | No | The maximum number of Lex nodes. |
| (Flows) ListFlowAliases requests per second | Each supported Region: 10 | No | The maximum number of ListFlowAliases requests per second. |
| (Flows) ListFlowVersions requests per second | Each supported Region: 10 | No | The maximum number of ListFlowVersions requests per second. |
| (Flows) ListFlows requests per second | Each supported Region: 10 | No | The maximum number of ListFlows requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Flows) Output nodes per flow | Each supported Region: 20 | No | The maximum number of flow output nodes. |
| (Flows) PrepareFlow requests per second | Each supported Region: 2 | No | The maximum number of PrepareFlow requests per second. |
| (Flows) Prompt nodes per flow | Each supported Region: 20 | Yes | The maximum number of prompt nodes. |
| (Flows) S3 retrieval nodes per flow | Each supported Region: 10 | No | The maximum number of S3 retrieval nodes. |
| (Flows) S3 storage nodes per flow | Each supported Region: 10 | No | The maximum number of S3 storage nodes. |
| (Flows) Total nodes per flow | Each supported Region: 40 | No | The maximum number of nodes in a flow. |
| (Flows) UpdateFlow requests per second | Each supported Region: 2 | No | The maximum number of UpdateFlow requests per second. |
| (Flows) UpdateFlowAlias requests per second | Each supported Region: 2 | No | The maximum number of UpdateFlowAlias requests per second. |
| (Flows) ValidateFlowDefinition requests per second | Each supported Region: 2 | No | The maximum number of ValidateFlowDefinition requests per second. |
| (Guardrails) Automated Reasoning policies per guardrail | Each supported Region: 2 | No | The maximum number of Automated Reasoning policies per guardrail. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Content policy maximum input size in text units (Classic tier) | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-northeast-2: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>eu-south-1: 25<br>eu-west-3: 25<br>sa-east-1: 25<br>Each of the other supported Regions: 106 | Yes | The maximum size of input text, measured in text units, that can be processed for content filters. While this limit applies to the classic tier, we recommend migrating to standard tier due to its superior robustness, additional capabilities, and multi-lingual support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Content policy maximum input size in text units (Standard tier - Recommended) | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 500<br>ap-northeast-2: 1,000<br>ap-south-1: 500<br>ap-southeast-1: 1,000<br>ap-southeast-2: 400<br>eu-central-1: 500<br>eu-south-1: 25<br>eu-west-1: 1,000<br>eu-west-3: 25<br>Each of the other supported Regions: 106 | Yes | The maximum size of input text, measured in text units, that can be processed for content filters. This applies to the standard tier, which is recommended. |
| (Guardrails) Contextual grounding policy maximum input size in text units | Each supported Region: 106 | Yes | The maximum size of input text, measured in text units, that can be processed by Contextual grounding policies |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Contextual grounding query length in text units | Each supported Region: 1 | No | The maximum length, in text units, of the query for contextual grounding |
| (Guardrails) Contextual grounding response length in text units | Each supported Region: 5 | No | The maximum length, in text units, of the response for contextual grounding |
| (Guardrails) Contextual grounding source length in text units | us-east-1: 100<br>us-west-2: 100<br>Each of the other supported Regions: 50 | No | The maximum length, in text units, of the grounding source for contextual grounding |
| (Guardrails) Example phrases per Topic | Each supported Region: 5 | No | The maximum number of topic examples that can be included per topic |
| (Guardrails) Guardrails per account | Each supported Region: 100 | No | The maximum number of guardrails in an account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Content filter policy text units burst rate (Classic tier) | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-northeast-2: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>Each of the other supported Regions: 25 | No | The maximum number of text units in one burst that can be processed for content filters. While this limit applies to the classic tier, we recommend migrating to standard tier due to its superior robustness, additional capabilities, and multi-lingual support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Content filter policy text units burst rate (Standard tier - Recommended) | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 500<br>ap-northeast-2: 1,000<br>ap-south-1: 500<br>ap-southeast-1: 1,000<br>ap-southeast-2: 400<br>eu-central-1: 500<br>eu-west-1: 1,000<br>Each of the other supported Regions: 25 | No | The maximum number of text units in one burst that can be processed for content filters. This applies to the standard tier, which is recommended. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Content filter policy text units per second (Classic tier) | us-east-1: 200<br><br>us-west-2: 200<br><br>Each of the other supported Regions: 25 | Yes | The maximum number of text units per second that can be processed for content filters. While this limit applies to the classic tier, we recommend migrating to standard tier due to its superior robustness, additional capabilities, and multi-lingual support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Content filter policy text units per second (Standard tier - Recommended) | us-east-1: 200<br>us-east-2: 200<br>us-west-1: 200<br>us-west-2: 200<br>ap-northeast-1: 50<br>ap-northeast-2: 100<br>ap-south-1: 50<br>ap-southeast-1: 100<br>eu-central-1: 50<br>eu-west-1: 200<br>Each of the other supported Regions: 25 | Yes | The maximum number of text units per second that can be processed for content filters. This applies to the standard tier, which is recommended. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Denied topic policy text units burst rate (Classic tier) | us-east-1: 200<br><br>us-west-2: 200<br><br>Each of the other supported Regions: 25 | No | The maximum number of text units in one burst that can be processed for denied topics. While this limit applies to the classic tier, we recommend migrating to standard tier due to its superior robustness, additional capabilities, and multi-lingual support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Denied topic policy text units burst rate (Standard tier - Recommended) | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 500<br>ap-northeast-2: 1,000<br>ap-south-1: 500<br>ap-southeast-1: 1,000<br>ap-southeast-2: 400<br>eu-central-1: 500<br>eu-west-1: 1,000<br>Each of the other supported Regions: 25 | No | The maximum number of text units in one burst that can be processed for denied topics. This applies to the standard tier, which is recommended. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Denied topic policy text units per second (Classic tier) | us-east-1: 50<br>us-west-2: 50<br>Each of the other supported Regions: 25 | Yes | The maximum number of text units per second that can be processed for denied topics. While this limit applies to the classic tier, we recommend migrating to standard tier due to its superior robustness, additional capabilities, and multi-lingual support. |
| (Guardrails) On-demand ApplyGuardrail Denied topic policy text units per second (Standard tier - Recommended) | us-east-1: 200<br>us-east-2: 75<br>us-west-1: 75<br>us-west-2: 200<br>eu-west-1: 200<br>Each of the other supported Regions: 25 | Yes | The maximum number of text units per second that can be processed for denied topics. This applies to the standard tier, which is recommended. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Sensitive information filter policy text units burst rate | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-northeast-2: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>Each of the other supported Regions: 25 | No | The maximum number of text units in one burst that can be processed for sensitive information filters. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Sensitive information filter policy text units per second | us-east-1: 1,000<br>us-east-2: 100<br>us-west-1: 50<br>us-west-2: 500<br>ap-northeast-1: 500<br>ap-northeast-2: 100<br>ap-northeast-3: 75<br>ap-south-1: 200<br>ap-south-2: 75<br>ap-southeast-1: 100<br>ap-southeast-2: 250<br>ca-central-1: 250<br>eu-central-1: 500<br>eu-central-2: 75<br>eu-north-1: 75<br>eu-south-1: 75<br>eu-south-2: 50 | Yes | The maximum number of text units per second that can be processed for sensitive information filters. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-west-1: 500<br>eu-west-2: 100<br>eu-west-3: 250<br>sa-east-1: 75<br>Each of the other supported Regions: 25 | | |
| (Guardrails) On-demand ApplyGuardrail Word filter policy text units burst rate | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-northeast-2: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>Each of the other supported Regions: 25 | No | The maximum number of text units in one burst that can be processed for word filters. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail Word filter policy text units per second | us-east-1: 500<br>us-east-2: 500<br>us-west-1: 500<br>us-west-2: 500<br>ap-northeast-1: 500<br>ap-northeast-2: 500<br>ap-south-1: 500<br>ap-southeast-1: 500<br>eu-central-1: 500<br>Each of the other supported Regions: 25 | Yes | The maximum number of text units per second that can be processed for word filters. |
| (Guardrails) On-demand ApplyGuardrail contextual grounding policy text units burst rate | Each supported Region: 106 | No | The maximum number of text units in one burst that can be processed for contextual grounding. |
| (Guardrails) On-demand ApplyGuardrail contextual grounding policy text units per second | Each supported Region: 106 | Yes | The maximum number of text units per second that can be processed for contextual grounding. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail requests burst rate | us-east-1: 100<br>us-east-2: 100<br>us-west-1: 100<br>us-west-2: 100<br>ap-northeast-1: 100<br>ap-northeast-2: 100<br>ap-south-1: 100<br>ap-southeast-1: 100<br>eu-central-1: 100<br>Each of the other supported Regions: 25 | No | The maximum number of ApplyGuardrail API calls that you can send in one burst. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) On-demand ApplyGuardrail requests per second | us-east-1: 100<br>us-east-2: 100<br>us-west-1: 100<br>us-west-2: 100<br>ap-northeast-1: 100<br>ap-northeast-2: 100<br>ap-south-1: 100<br>ap-southeast-1: 100<br>eu-central-1: 100<br>Each of the other supported Regions: 25 | Yes | The maximum number of ApplyGuardrail API calls allowed per second |
| (Guardrails) On-demand InvokeGuardrailChecks requests burst rate | Each supported Region: 1,500 | No | The maximum number of InvokeGuardrailChecks API calls that you can send in one burst |
| (Guardrails) On-demand InvokeGuardrailChecks requests per minute | Each supported Region: 1,500 | Yes | The maximum number of InvokeGuardrailChecks API calls allowed per minute |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Regex entities in Sensitive Information Filter | Each supported Region: 30 | No | The maximum number of guardrail filter regexes that can be included in a sensitive information policy |
| (Guardrails) Regex length in characters | Each supported Region: 500 | No | The maximum length, in characters, of a guardrail filter regex |
| (Guardrails) Sensitive information policy maximum input size in text units | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-northeast-2: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>Each of the other supported Regions: 106 | Yes | The maximum size of input text, measured in text units, that can be processed by Sensitive information filter policies |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Topic policy maximum input size in text units (Classic tier) | us-east-1: 200<br><br>us-west-2: 200<br><br>ap-southeast-1: 25<br><br>eu-south-1: 25<br><br>eu-west-3: 25<br><br>sa-east-1: 25<br><br>Each of the other supported Regions: 106 | Yes | The maximum size of input text, measured in text units, that can be processed for denied topics. While this limit applies to the classic tier, we recommend migrating to standard tier due to its superior robustness, additional capabilities, and multi-lingual support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Topic policy maximum input size in text units (Standard tier - Recommended) | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 500<br>ap-northeast-2: 1,000<br>ap-south-1: 500<br>ap-southeast-1: 1,000<br>ap-southeast-2: 400<br>eu-central-1: 500<br>eu-south-1: 25<br>eu-west-1: 1,000<br>eu-west-3: 25<br>Each of the other supported Regions: 106 | Yes | The maximum size of input text, measured in text units, that can be processed for denied topics. This applies to the standard tier, which is recommended. |
| (Guardrails) Topics per guardrail | Each supported Region: 30 | No | The maximum number of topics that can be defined across guardrail topic policies |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Guardrails) Versions per guardrail | Each supported Region: 20 | No | The maximum number of versions that a guardrail can have |
| (Guardrails) Word length in characters | Each supported Region: 100 | No | The maximum length of a word, in characters, in a blocked word list |
| (Guardrails) Word policy maximum input size in text units | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-northeast-2: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>Each of the other supported Regions: 106 | Yes | The maximum size of input text, measured in text units, that can be processed by Word filter policies |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Guardrails) Words per word policy | Each supported Region: 10,000 | No | The maximum number of words that can be included in a blocked word list |
| (Knowledge Bases) Concurrent IngestKnowledgeBaseDocuments and DeleteKnowledgeBaseDocuments requests per account | Each supported Region: 10 | No | The maximum number of IngestKnowledgeBaseDocuments and DeleteKnowledgeBaseDocuments requests that can be running at the same time in an account. |
| (Knowledge Bases) Concurrent ingestion jobs per account | Each supported Region: 5 | No | The maximum number of ingestion jobs that can be running at the same time in an account. |
| (Knowledge Bases) Concurrent ingestion jobs per data source | Each supported Region: 1 | No | The maximum number of ingestion jobs that can be running at the same time for a data source. |
| (Knowledge Bases) Concurrent ingestion jobs per knowledge base | Each supported Region: 1 | No | The maximum number of ingestion jobs that can be running at the same time for a knowledge base. |
| (Knowledge Bases) CreateDataSource requests per second | Each supported Region: 2 | No | The maximum number of CreateDataSource API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Knowledge Bases) CreateKnowledgeBase requests per second | Each supported Region: 2 | No | The maximum number of CreateKnowledgeBase API requests per second. |
| (Knowledge Bases) Data sources per knowledge base | Each supported Region: 5 | No | The maximum number of data sources per knowledge base. |
| (Knowledge Bases) DeleteDataSource requests per second | Each supported Region: 2 | No | The maximum number of DeleteDataSource API requests per second. |
| (Knowledge Bases) DeleteKnowledgeBase requests per second | Each supported Region: 2 | No | The maximum number of DeleteKnowledgeBase API requests per second. |
| (Knowledge Bases) DeleteKnowledgeBaseDocuments requests per second | Each supported Region: 5 | No | The maximum number of DeleteKnowledgeBaseDocuments API requests per second. |
| (Knowledge Bases) Files to add or update per ingestion job | Each supported Region: 5,000,000 | No | The maximum number of new and updated files that can be ingested per ingestion job. |
| (Knowledge Bases) Files to delete per ingestion job | Each supported Region: 5,000,000 | No | The maximum number of files that can be deleted per ingestion job. |
| (Knowledge Bases) Files to ingest per IngestKnowledgeBaseDocuments job. | Each supported Region: 25 | No | The maximum number of documents that can be ingested per IngestKnowledgeBaseDocuments request. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Knowledge Bases) GenerateQuery requests per second | Each supported Region: 2 | No | The maximum number of GenerateQuery API requests per second. |
| (Knowledge Bases) GetDataSource requests per second | Each supported Region: 10 | No | The maximum number of GetDataSource API requests per second. |
| (Knowledge Bases) GetIngestionJob requests per second | Each supported Region: 10 | No | The maximum number of GetIngestionJob API requests per second. |
| (Knowledge Bases) GetKnowledgeBase requests per second | Each supported Region: 10 | No | The maximum number of GetKnowledgeBase API requests per second. |
| (Knowledge Bases) GetKnowledgeBaseDocuments requests per second | Each supported Region: 5 | No | The maximum number of GetKnowledgeBaseDocuments API requests per second. |
| (Knowledge Bases) IngestKnowledgeBaseDocuments requests per second | Each supported Region: 5 | No | The maximum number of IngestKnowledgeBaseDocuments API requests per second. |
| (Knowledge Bases) IngestKnowledgeBaseDocuments total payload size | Each supported Region: 6 | No | The maximum size (in MB) of total payload in an IngestKnowledgeBaseDocuments request. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Knowledge Bases) Ingestion job file size with text content | Each supported Region: 50 | No | The maximum size (in MB) of a file with text content (such as .txt, .pdf, or .docx files) in an ingestion job. |
| (Knowledge Bases) Ingestion job size | Each supported Region: 100 | No | The maximum size (in GB) of an ingestion job. |
| (Knowledge Bases) Knowledge bases per account | Each supported Region: 100 | No | The maximum number of knowledge bases per account. |
| (Knowledge Bases) ListDataSources requests per second | Each supported Region: 10 | No | The maximum number of ListDataSources API requests per second. |
| (Knowledge Bases) ListIngestionJobs requests per second | Each supported Region: 10 | No | The maximum number of ListIngestionJobs API requests per second. |
| (Knowledge Bases) ListKnowledgeBaseDocuments requests per second | Each supported Region: 5 | No | The maximum number of ListKnowledgeBaseDocuments API requests per second. |
| (Knowledge Bases) ListKnowledgeBases requests per second | Each supported Region: 10 | No | The maximum number of ListKnowledgeBases API requests per second. |
| (Knowledge Bases) Maximum number of files for BDA parser | Each supported Region: 1,000 | No | The maximum number of files that can be used with Amazon Bedrock Data Automation as parser. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Knowledge Bases) Maximum number of files for Foundation Models as a parser | Each supported Region: 1,000 | No | The maximum number of files that can be used with Foundation Models as a parser. |
| (Knowledge Bases) Rerank requests per second | Each supported Region: 10 | No | The maximum number of Rerank API requests per second. |
| (Knowledge Bases) Retrieve requests per second | Each supported Region: 20 | No | The maximum number of Retrieve API requests per second. |
| (Knowledge Bases) RetrieveAndGenerate requests per second | Each supported Region: 20 | No | The maximum number of RetrieveAndGenerate API requests per second. |
| (Knowledge Bases) RetrieveAndGenerateStream requests per second | Each supported Region: 20 | No | The maximum number of RetrieveAndGenerateStream API requests per second. |
| (Knowledge Bases) StartIngestionJob requests per second | Each supported Region: 0.1 | No | The maximum number of StartIngestionJob API requests per second. |
| (Knowledge Bases) UpdateDataSource requests per second | Each supported Region: 2 | No | The maximum number of UpdateDataSource API requests per second. |
| (Knowledge Bases) UpdateKnowledgeBase requests per second | Each supported Region: 2 | No | The maximum number of UpdateKnowledgeBase API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Knowledge Bases) User query size | Each supported Region: 1,000 | No | The maximum size (in characters) of a user query. |
| (Managed Knowledge Bases) AgenticRetrieveStream requests per minute per account | Each supported Region: 60 | No | The maximum number of AgenticRetrieveStream API requests per minute per account for Managed KBs. |
| (Managed Knowledge Bases) AgenticRetrieveStream user query size | Each supported Region: 10,000 | No | The maximum size (in characters) of a user query for AgenticRetrieveStream for Managed KBs. |
| (Managed Knowledge Bases) Concurrent ingestion jobs per knowledge base | Each supported Region: 50 | No | The maximum number of concurrent ingestion jobs per Managed KB. |
| (Managed Knowledge Bases) Data sources per knowledge base | Each supported Region: 200 | No | The maximum number of data sources per Managed KB. |
| (Managed Knowledge Bases) DeleteKnowledgeBaseDocuments requests per second | Each supported Region: 10 | No | The maximum number of DeleteKnowledgeBaseDocuments API requests per second for Managed KBs. |
| (Managed Knowledge Bases) DeleteResourcePolicy requests per second | Each supported Region: 5 | No | The maximum number of DeleteResourcePolicy API requests per second for Managed KBs. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Managed Knowledge Bases) Files to ingest per IngestKnowledgeBaseDocuments request | Each supported Region: 10 | No | The maximum number of files to ingest per IngestKnowledgeBaseDocuments API request for Managed KBs. |
| (Managed Knowledge Bases) GetDocumentContent requests per second per account | Each supported Region: 100 | No | The maximum number of GetDocumentContent API requests per second per account. |
| (Managed Knowledge Bases) GetDocumentContent requests per second per knowledge base | Each supported Region: 5 | No | The maximum number of GetDocumentContent API requests per second per Managed KB. |
| (Managed Knowledge Bases) GetResourcePolicy requests per second | Each supported Region: 5 | No | The maximum number of GetResourcePolicy API requests per second for Managed KBs. |
| (Managed Knowledge Bases) Individual file extracted text size (MB) | Each supported Region: 30 | No | The maximum size (in MB) of extracted text from a single file for Managed KBs. |
| (Managed Knowledge Bases) IngestKnowledgeBaseDocuments requests per second | Each supported Region: 20 | No | The maximum number of IngestKnowledgeBaseDocuments API requests per second for Managed KBs. |
| (Managed Knowledge Bases) Knowledge bases per account | Each supported Region: 10,000 | No | The maximum number of Managed KBs per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Managed Knowledge Bases) ListKnowledgeBaseDocuments requests per second | Each supported Region: 10 | No | The maximum number of ListKnowledgeBaseDocuments API requests per second for Managed KBs. |
| (Managed Knowledge Bases) PutResourcePolicy requests per second | Each supported Region: 5 | No | The maximum number of PutResourcePolicy API requests per second for Managed KBs. |
| (Managed Knowledge Bases) Retrieve requests per minute per knowledge base | Each supported Region: 600 | No | The maximum number of Retrieve API requests per minute per Managed KB. |
| (Managed Knowledge Bases) Retrieve requests per second per account | Each supported Region: 100 | No | The maximum number of Retrieve API requests per second per account for Managed KBs. |
| (Managed Knowledge Bases) Retrieve user query size | Each supported Region: 10,000 | No | The maximum size (in characters) of a user query for Retrieve for Managed KBs. |
| (Managed Knowledge Bases) Total storage size per knowledge base (TB) | Each supported Region: 10 | No | The maximum total storage size (in TB) per Managed KB. |
| (Model customization) Custom models per account | Each supported Region: 100 | Yes | The maximum number of custom models in an account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Model customization) In-progress custom model deployments | Each supported Region: 2 | Yes | The maximum number of in-progress custom model deployments |
| (Model customization) Maximum input file size for distillation customization jobs | Each supported Region: 2 Gigabytes | No | The maximum input file size for distillation customization jobs. |
| (Model customization) Maximum line length for distillation customization jobs | Each supported Region: 16 Kilobytes | No | The maximum line length in input file for distillation customization jobs. |
| (Model customization) Maximum number of prompts for distillation customization jobs | Each supported Region: 15,000 | No | The maximum number of prompts required for distillation customization jobs. |
| (Model customization) Maximum number of training records for an Amazon Nova Canvas Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum number of records allowed for an Amazon Nova Canvas Fine-tuning job. |
| (Model customization) Maximum student model fine tuning context length for Amazon Nova Micro V1 distillation customization jobs | Each supported Region: 32,000 | No | The maximum student model fine tuning context length for Amazon Nova Micro V1 distillation customization jobs. |
| (Model customization) Maximum student model fine tuning context length for Amazon Nova V1 distillation customization jobs | Each supported Region: 32,000 | No | The maximum student model fine tuning context length for Amazon Nova V1 distillation customization jobs. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Model customization) Maximum student model fine tuning context length for Anthropic Claude 3 haiku 20240307 V1 distillation customization jobs | Each supported Region: 32,000 | No | The maximum student model fine tuning context length for Anthropic Claude 3 haiku 20240307 V1 distillation customization jobs. |
| (Model customization) Maximum student model fine tuning context length for Llama 3.1 70B Instruct V1 distillation customization jobs | Each supported Region: 16,000 | No | The maximum student model fine tuning context length for Llama 3.1 70B Instruct V1 distillation customization jobs. |
| (Model customization) Maximum student model fine tuning context length for Llama 3.1 8B Instruct V1 distillation customization jobs | Each supported Region: 32,000 | No | The maximum student model fine tuning context length for Llama 3.1 8B Instruct V1 distillation customization jobs. |
| (Model customization) Minimum number of prompts for distillation customization jobs | Each supported Region: 100 | No | The minimum number of prompts required for distillation customization jobs. |
| (Model customization) Scheduled customization jobs | Each supported Region: 10 | No | The maximum number of scheduled customization jobs. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of on demand custom model deployment requests per minute for Amazon Nova 2 Lite | Each supported Region: 2,000 | No | The sum of input and output on-demand custom model deployment requests per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova 2 Lite |
| (Model customization) Sum of on demand custom model deployment requests per minute for Amazon Nova Lite | Each supported Region: 2,000 | No | The sum of input and output on-demand custom model deployment requests per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Lite |
| (Model customization) Sum of on demand custom model deployment requests per minute for Amazon Nova Micro | Each supported Region: 2,000 | No | The sum of input and output on-demand custom model deployment requests per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Micro |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of on demand custom model deployment requests per minute for Amazon Nova Pro | Each supported Region: 200 | No | The sum of input and output on-demand custom model deployment requests per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Pro |
| (Model customization) Sum of on demand custom model deployment requests per minute for Meta Llama 3.3 70B Instruct | Each supported Region: 400 | No | The sum of input and output on-demand custom model deployment requests per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Meta Llama 3.3 70B Instruct |
| (Model customization) Sum of on demand custom model deployment tokens per day for Amazon Nova 2 Lite | Each supported Region: 5,760,000,000 | No | The sum of input and output on-demand custom model deployment tokens per day submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova 2 Lite |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Model customization) Sum of on demand custom model deployment tokens per day for Amazon Nova Lite | Each supported Region: 5,760,000,000 | No | The sum of input and output on-demand custom model deployment tokens per day submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Lite |
| (Model customization) Sum of on demand custom model deployment tokens per day for Amazon Nova Micro | Each supported Region: 5,760,000,000 | No | The sum of input and output on-demand custom model deployment tokens per day submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Micro |
| (Model customization) Sum of on demand custom model deployment tokens per day for Amazon Nova Pro | Each supported Region: 1,152,000,000 | No | The sum of input and output on-demand custom model deployment tokens per day submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Pro |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of on demand custom model deployment tokens per day for Meta Llama 3.3 70B Instruct | Each supported Region: 432,000,000 | No | The sum of input and output on-demand custom model deployment tokens per day submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Meta Llama 3.3 70B Instruct |
| (Model customization) Sum of on demand custom model deployment tokens per minute for Amazon Nova 2 Lite | Each supported Region: 4,000,000 | No | The sum of input and output on-demand custom model deployment tokens per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova 2 Lite |
| (Model customization) Sum of on demand custom model deployment tokens per minute for Amazon Nova Lite | Each supported Region: 4,000,000 | No | The sum of input and output on-demand custom model deployment tokens per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Lite |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of on demand custom model deployment tokens per minute for Amazon Nova Micro | Each supported Region: 4,000,000 | No | The sum of input and output on-demand custom model deployment tokens per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Micro |
| (Model customization) Sum of on demand custom model deployment tokens per minute for Amazon Nova Pro | Each supported Region: 800,000 | No | The sum of input and output on-demand custom model deployment tokens per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Amazon Nova Pro |
| (Model customization) Sum of on demand custom model deployment tokens per minute for Meta Llama 3.3 70B Instruct | Each supported Region: 300,000 | No | The sum of input and output on-demand custom model deployment tokens per minute submitted to the Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream actions for Meta Llama 3.3 70B Instruct |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of training and validation records for a Amazon Nova 2 Lite Fine-tuning job | Each supported Region: 20,000 | Yes | The maximum combined number of training and validation records allowed for a Amazon Nova 2 Lite Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Amazon Nova Lite Fine-tuning job | Each supported Region: 20,000 | Yes | The maximum combined number of training and validation records allowed for a Amazon Nova Lite Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Amazon Nova Micro Fine-tuning job | Each supported Region: 20,000 | Yes | The maximum combined number of training and validation records allowed for a Amazon Nova Micro Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Amazon Nova Pro Fine-tuning job | Each supported Region: 20,000 | Yes | The maximum combined number of training and validation records allowed for a Amazon Nova Pro Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Claude 3 Haiku v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Claude 3 Haiku Fine-tuning job. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of training and validation records for a Claude 3-5-Haiku v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Claude 3-5-Haiku Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 2 13B v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 2 13B Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 2 70B v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 2 70B Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 3.1 70B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.1 70B Instruct Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 3.1 8B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.1 8B Instruct Fine-tuning job. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of training and validation records for a Meta Llama 3.2 11B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.2 11B Instruct Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 3.2 1B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.2 1B Instruct Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 3.2 3B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.2 3B Instruct Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 3.2 90B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.2 90B Instruct Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Meta Llama 3.3 70B Instruct v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Meta Llama 3.3 70B Instruct Fine-tuning job. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Model customization) Sum of training and validation records for a Titan Image Generator G1 V1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Image Generator Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Titan Image Generator G1 V2 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Image Generator V2 Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Titan Multimodal Embeddings G1 v1 Fine-tuning job | Each supported Region: 50,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Multimodal Embeddings Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Titan Text G1 - Express v1 Continued Pre-Training job | Each supported Region: 100,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Text Express Continued Pre-Training job. |
| (Model customization) Sum of training and validation records for a Titan Text G1 - Express v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Text Express Fine-tuning job. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| (Model customization) Sum of training and validation records for a Titan Text G1 - Lite v1 Continued Pre-Training job | Each supported Region: 100,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Text Lite Continued Pre-Training job. |
| (Model customization) Sum of training and validation records for a Titan Text G1 - Lite v1 Fine-tuning job | Each supported Region: 10,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Text Lite Fine-tuning job. |
| (Model customization) Sum of training and validation records for a Titan Text G1 - Premier v1 Fine-tuning job | Each supported Region: 20,000 | Yes | The maximum combined number of training and validation records allowed for a Titan Text Premier Fine-tuning job. |
| (Model customization) Total number of custom model deployments | Each supported Region: 10 | Yes | Total number of custom model deployments |
| (Prompt management) CreatePrompt requests per second | Each supported Region: 2 | No | The maximum number of CreatePrompt requests per second. |
| (Prompt management) CreatePromptVersion requests per second | Each supported Region: 2 | No | The maximum number of CreatePromptVersion requests per second. |
| (Prompt management) DeletePrompt requests per second | Each supported Region: 2 | No | The maximum number of DeletePrompt requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| (Prompt management) GetPrompt requests per second | Each supported Region: 10 | No | The maximum number of GetPrompt requests per second. |
| (Prompt management) ListPrompts requests per second | Each supported Region: 10 | No | The maximum number of ListPrompts requests per second. |
| (Prompt management) Prompts per account | Each supported Region: 500 | Yes | The maximum number of prompts. |
| (Prompt management) UpdatePrompt requests per second | Each supported Region: 2 | No | The maximum number of UpdatePrompt requests per second. |
| (Prompt management) Versions per prompt | Each supported Region: 10 | No | The maximum number of versions per prompt. |
| APIs per Agent | Each supported Region: 11 | Yes | The maximum number of APIs that you can add to an Agent. |
| Action groups per Agent | Each supported Region: 20 | Yes | The maximum number of action groups that you can add to an Agent. |
| Agent Collaborators per Agent | Each supported Region: 1,000 | Yes | The maximum number of collaborator agents that you can add to an Agent. |
| Agents per account | Each supported Region: 1,000 | Yes | The maximum number of Agents in one account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AssociateAgentKnowledgeBase requests per second | Each supported Region: 6 | No | The maximum number of AssociateAgentKnowledgeBase API requests per second. |
| Associated aliases per Agent | Each supported Region: 10 | No | The maximum number of aliases that you can associate with an Agent. |
| Associated knowledge bases per Agent | Each supported Region: 2 | Yes | The maximum number of knowledge bases that you can associate with an Agent. |
| Batch inference input file size (in GB) for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Amazon Nova 2 Multimodal Embeddings V1. |
| Batch inference input file size (in GB) for Amazon Nova Premier | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Amazon Nova Premier. |
| Batch inference input file size (in GB) for Claude 3 Haiku | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude 3 Haiku. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Claude 3 Opus | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude 3 Opus. |
| Batch inference input file size (in GB) for Claude 3 Sonnet | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude 3 Sonnet. |
| Batch inference input file size (in GB) for Claude 3.5 Haiku | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude 3.5 Haiku. |
| Batch inference input file size (in GB) for Claude 3.5 Sonnet | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude 3.5 Sonnet. |
| Batch inference input file size (in GB) for Claude 3.5 Sonnet v2 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude 3.5 Sonnet v2. |
| Batch inference input file size (in GB) for Claude 3.7 Sonnet | Each supported Region: 1 | Yes | The maximum size of a single file (in GB) submitted for batch inference for Claude 3.7 Sonnet. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Claude Haiku 4.5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude Haiku 4.5. |
| Batch inference input file size (in GB) for Claude Opus 4.5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude Opus 4.5. |
| Batch inference input file size (in GB) for Claude Opus 4.6 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude Opus 4.6. |
| Batch inference input file size (in GB) for Claude Opus 5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude Opus 5. |
| Batch inference input file size (in GB) for Claude Sonnet 4 | Each supported Region: 1 | Yes | The maximum size of a single file (in GB) submitted for batch inference for Claude Sonnet 4. |
| Batch inference input file size (in GB) for Claude Sonnet 4.5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude Sonnet 4.5. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Claude Sonnet 4.6 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Claude Sonnet 4.6. |
| Batch inference input file size (in GB) for DeepSeek V3.2 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for DeepSeek V3.2. |
| Batch inference input file size (in GB) for DeepSeek v3 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for DeepSeek v3. |
| Batch inference input file size (in GB) for Devstral 2 123B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Devstral 2 123B. |
| Batch inference input file size (in GB) for GLM 4.7 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for GLM 4.7. |
| Batch inference input file size (in GB) for GLM 4.7 Flash | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for GLM 4.7 Flash. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for GLM 5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for GLM 5. |
| Batch inference input file size (in GB) for Gemma 3 12B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Gemma 3 12B. |
| Batch inference input file size (in GB) for Gemma 3 27B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Gemma 3 27B. |
| Batch inference input file size (in GB) for Gemma 3 4B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Gemma 3 4B. |
| Batch inference input file size (in GB) for Kimi K2 Thinking | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Kimi K2 Thinking. |
| Batch inference input file size (in GB) for Kimi K2.5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Kimi K2.5. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Llama 3.1 405B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.1 405B Instruct. |
| Batch inference input file size (in GB) for Llama 3.1 70B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.1 70B Instruct. |
| Batch inference input file size (in GB) for Llama 3.1 8B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.1 8B Instruct. |
| Batch inference input file size (in GB) for Llama 3.2 11B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.2 11B Instruct. |
| Batch inference input file size (in GB) for Llama 3.2 1B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference Llama 3.2 1B Instruct. |
| Batch inference input file size (in GB) for Llama 3.2 3B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.2 3B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Llama 3.2 90B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.2 90B Instruct. |
| Batch inference input file size (in GB) for Llama 3.3 70B Instruct | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 3.3 70B Instruct. |
| Batch inference input file size (in GB) for Llama 4 Maverick | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 4 Maverick. |
| Batch inference input file size (in GB) for Llama 4 Scout | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Llama 4 Scout. |
| Batch inference input file size (in GB) for Magistral Small 2509 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Magistral Small 2509. |
| Batch inference input file size (in GB) for MiniMax M2 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for MiniMax M2. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for MiniMax M2.1 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for MiniMax M2.1. |
| Batch inference input file size (in GB) for MiniMax M2.5 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for MiniMax M2.5. |
| Batch inference input file size (in GB) for Ministral 3 14B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Ministral 3 14B. |
| Batch inference input file size (in GB) for Ministral 3 8B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Ministral 3 8B. |
| Batch inference input file size (in GB) for Ministral 3B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Ministral 3B. |
| Batch inference input file size (in GB) for Mistral Large 2 (24.07) | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Mistral Large 2 (24.07). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Mistral Large 3 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Mistral Large 3. |
| Batch inference input file size (in GB) for Mistral Small | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Mistral Small. |
| Batch inference input file size (in GB) for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for NVIDIA Nemotron 3 Super 120B A12B. |
| Batch inference input file size (in GB) for NVIDIA Nemotron Nano 12B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for NVIDIA Nemotron Nano 12B. |
| Batch inference input file size (in GB) for NVIDIA Nemotron Nano 3 30B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for NVIDIA Nemotron Nano 3 30B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for NVIDIA Nemotron Nano 9B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for NVIDIA Nemotron Nano 9B. |
| Batch inference input file size (in GB) for Nova 2 Lite | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Nova 2 Lite. |
| Batch inference input file size (in GB) for Nova Lite V1 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Nova Lite V1. |
| Batch inference input file size (in GB) for Nova Micro V1 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Nova Micro V1. |
| Batch inference input file size (in GB) for Nova Pro V1 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Nova Pro V1. |
| Batch inference input file size (in GB) for OpenAI GPT OSS 120b | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for OpenAI GPT OSS 120b. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for OpenAI GPT OSS 20b | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for OpenAI GPT OSS 20b. |
| Batch inference input file size (in GB) for OpenAI GPT OSS Safeguard 120b | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for OpenAI GPT OSS Safeguard 120b. |
| Batch inference input file size (in GB) for OpenAI GPT OSS Safeguard 20b | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for OpenAI GPT OSS Safeguard 20b. |
| Batch inference input file size (in GB) for Qwen3 235B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 235B. |
| Batch inference input file size (in GB) for Qwen3 32B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 32B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Qwen3 Coder 30B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 Coder 30B. |
| Batch inference input file size (in GB) for Qwen3 Coder 480B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 Coder 480B. |
| Batch inference input file size (in GB) for Qwen3 Coder Next | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 Coder Next. |
| Batch inference input file size (in GB) for Qwen3 Next 80B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 Next 80B. |
| Batch inference input file size (in GB) for Qwen3 VL 235B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Qwen3 VL 235B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference input file size (in GB) for Titan Multimodal Embeddings G1 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Titan Multimodal Embeddings G1. |
| Batch inference input file size (in GB) for Titan Text Embeddings V2 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Titan Text Embeddings V2. |
| Batch inference input file size (in GB) for Voxtral Mini 3B 2507 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Voxtral Mini 3B 2507. |
| Batch inference input file size (in GB) for Voxtral Small 24B 2507 | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Voxtral Small 24B 2507. |
| Batch inference input file size (in GB) for Writer Palmyra Vision 7B | Each supported Region: 1 | No | The maximum size of a single file (in GB) submitted for batch inference for Writer Palmyra Vision 7B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Qwen3 Next 80B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 Next 80B. |
| Batch inference job size (in GB) for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 100 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Amazon Nova 2 Multimodal Embeddings V1. |
| Batch inference job size (in GB) for Amazon Nova Premier | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Amazon Nova Premier. |
| Batch inference job size (in GB) for Claude 3 Haiku | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3 Haiku. |
| Batch inference job size (in GB) for Claude 3 Opus | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3 Opus. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Claude 3 Sonnet | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3 Sonnet. |
| Batch inference job size (in GB) for Claude 3.5 Haiku | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3.5 Haiku. |
| Batch inference job size (in GB) for Claude 3.5 Sonnet | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3.5 Sonnet. |
| Batch inference job size (in GB) for Claude 3.5 Sonnet v2 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3.5 Sonnet v2. |
| Batch inference job size (in GB) for Claude 3.7 Sonnet | Each supported Region: 5 | Yes | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude 3.7 Sonnet. |
| Batch inference job size (in GB) for Claude Haiku 4.5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Haiku 4.5. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Claude Opus 4.5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Opus 4.5. |
| Batch inference job size (in GB) for Claude Opus 4.6 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Opus 4.6. |
| Batch inference job size (in GB) for Claude Opus 5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Opus 5. |
| Batch inference job size (in GB) for Claude Sonnet 4 | Each supported Region: 5 | Yes | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Sonnet 4. |
| Batch inference job size (in GB) for Claude Sonnet 4.5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Sonnet 4.5. |
| Batch inference job size (in GB) for Claude Sonnet 4.6 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Claude Sonnet 4.6. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for DeepSeek V3.2 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for DeepSeek V3.2. |
| Batch inference job size (in GB) for DeepSeek v3 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for DeepSeek v3. |
| Batch inference job size (in GB) for Devstral 2 123B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Devstral 2 123B. |
| Batch inference job size (in GB) for GLM 4.7 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for GLM 4.7. |
| Batch inference job size (in GB) for GLM 4.7 Flash | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for GLM 4.7 Flash. |
| Batch inference job size (in GB) for GLM 5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for GLM 5. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Gemma 3 12B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Gemma 3 12B. |
| Batch inference job size (in GB) for Gemma 3 27B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Gemma 3 27B. |
| Batch inference job size (in GB) for Gemma 3 4B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Gemma 3 4B. |
| Batch inference job size (in GB) for Kimi K2 Thinking | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Kimi K2 Thinking. |
| Batch inference job size (in GB) for Kimi K2.5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Kimi K2.5. |
| Batch inference job size (in GB) for Llama 3.1 405B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.1 405B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Llama 3.1 70B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.1 70B Instruct. |
| Batch inference job size (in GB) for Llama 3.1 8B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.1 8B Instruct. |
| Batch inference job size (in GB) for Llama 3.2 11B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.2 11B Instruct. |
| Batch inference job size (in GB) for Llama 3.2 1B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.2 1B Instruct. |
| Batch inference job size (in GB) for Llama 3.2 3B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.2 3B Instruct. |
| Batch inference job size (in GB) for Llama 3.2 90B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.2 90B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Llama 3.3 70B Instruct | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 3.3 70B Instruct. |
| Batch inference job size (in GB) for Llama 4 Maverick | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 4 Maverick. |
| Batch inference job size (in GB) for Llama 4 Scout | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Llama 4 Scout. |
| Batch inference job size (in GB) for Magistral Small 2509 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Magistral Small 2509. |
| Batch inference job size (in GB) for MiniMax M2 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for MiniMax M2. |
| Batch inference job size (in GB) for MiniMax M2.1 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for MiniMax M2.1. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for MiniMax M2.5 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for MiniMax M2.5. |
| Batch inference job size (in GB) for Ministral 3 14B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Ministral 3 14B. |
| Batch inference job size (in GB) for Ministral 3 8B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Ministral 3 8B. |
| Batch inference job size (in GB) for Ministral 3B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Ministral 3B. |
| Batch inference job size (in GB) for Mistral Large 2 (24.07) | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Mistral Large 2 (24.07). |
| Batch inference job size (in GB) for Mistral Large 3 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Mistral Large 3. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Batch inference job size (in GB) for Mistral Small | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Mistral Small. |
| Batch inference job size (in GB) for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for NVIDIA Nemotron 3 Super 120B A12B. |
| Batch inference job size (in GB) for NVIDIA Nemotron Nano 12B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for NVIDIA Nemotron Nano 12B. |
| Batch inference job size (in GB) for NVIDIA Nemotron Nano 3 30B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for NVIDIA Nemotron Nano 3 30B. |
| Batch inference job size (in GB) for NVIDIA Nemotron Nano 9B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for NVIDIA Nemotron Nano 9B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Nova Lite V1 | Each supported Region: 100 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Nova Lite V1. |
| Batch inference job size (in GB) for Nova Micro V1 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Nova Micro V1. |
| Batch inference job size (in GB) for Nova Pro V1 | Each supported Region: 100 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Nova Pro V1. |
| Batch inference job size (in GB) for OpenAI GPT OSS 120b | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for OpenAI GPT OSS 120b. |
| Batch inference job size (in GB) for OpenAI GPT OSS 20b | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for OpenAI GPT OSS 20b. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for OpenAI GPT OSS Safeguard 120b | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for OpenAI GPT OSS Safeguard 120b. |
| Batch inference job size (in GB) for OpenAI GPT OSS Safeguard 20b | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for OpenAI GPT OSS Safeguard 20b. |
| Batch inference job size (in GB) for Qwen3 235B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 235B. |
| Batch inference job size (in GB) for Qwen3 32B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 32B. |
| Batch inference job size (in GB) for Qwen3 Coder 30B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 Coder 30B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Qwen3 Coder 480B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 Coder 480B. |
| Batch inference job size (in GB) for Qwen3 Coder Next | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 Coder Next. |
| Batch inference job size (in GB) for Qwen3 VL 235B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Qwen3 VL 235B. |
| Batch inference job size (in GB) for Titan Multimodal Embeddings G1 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Titan Multimodal Embeddings G1. |
| Batch inference job size (in GB) for Titan Text Embeddings V2 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Titan Text Embeddings V2. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch inference job size (in GB) for Voxtral Mini 3B 2507 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Voxtral Mini 3B 2507. |
| Batch inference job size (in GB) for Voxtral Small 24B 2507 | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Voxtral Small 24B 2507. |
| Batch inference job size (in GB) for Writer Palmyra Vision 7B | Each supported Region: 5 | No | The maximum cumulative size of all input files (in GB) included in the batch inference job for Writer Palmyra Vision 7B. |
| Characters in Agent instructions | Each supported Region: 20,000 | No | The maximum number of characters in the instructions for an Agent. |
| Concurrent model import jobs | Each supported Region: 1 | No | The maximum number of model import jobs that are concurrently in progress. |
| CreateAgent requests per second | Each supported Region: 6 | No | The maximum number of CreateAgent API requests per second. |
| CreateAgentActionGroup requests per second | Each supported Region: 12 | No | The maximum number of CreateAgentActionGroup API requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateAgentAlias requests per second | Each supported Region: 2 | No | The maximum number of CreateAgentAlias API requests per second. |
| Cross-Region model inference requests per minute for Anthropic Claude 3.5 Haiku | Each supported Region: 2,000 | No | The maximum number of times that you can call InvokeModel and InvokeModelWithResponseStream in one minute. The quota considers the combined sum of InvokeModel and InvokeModelWithResponseStream tokens for Anthropic Claude 3.5 Haiku. |
| Cross-Region model inference requests per minute for Anthropic Claude 3.5 Sonnet V2 | Each supported Region: 100 | No | The maximum number of times that you can call model inference in one minute for Anthropic Claude 3.5 Sonnet V2. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-Region model inference tokens per minute for Anthropic Claude 3.5 Haiku | Each supported Region: 4,000,000 | Yes | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3.5 Haiku. |
| Cross-Region model inference tokens per minute for Anthropic Claude 3.5 Sonnet V2 | Each supported Region: 800,000 | Yes | The maximum number of tokens that you can submit for model inference in one minute for Anthropic Claude 3.5 Sonnet V2. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Amazon Nova 2 Lite | Each supported Region: 2,000 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova 2 Lite. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Amazon Nova 2 Omni | Each supported Region: 2,000 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova 2 Omni. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Amazon Nova 2 Pro Preview | Each supported Region: 100 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova 2 Pro Preview. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Amazon Nova Lite | us-east-1: 4,000<br>us-east-2: 4,000<br>us-west-1: 4,000<br>us-west-2: 4,000<br>Each of the other supported Regions: 400 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova Lite. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Amazon Nova Micro | us-east-1: 4,000<br>us-east-2: 4,000<br>us-west-2: 4,000<br>Each of the other supported Regions: 400 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova Micro. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Amazon Nova Premier V1 | Each supported Region: 500 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova Premier V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Amazon Nova Pro | Each supported Region: 500 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Amazon Nova Pro. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude 3 Haiku | us-east-1: 2,000<br>us-west-2: 2,000<br>ap-northeast-1: 400<br>ap-southeast-1: 400<br>Each of the other supported Regions: 800 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Haiku. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude 3 Opus | Each supported Region: 100 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude 3 Opus. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude 3 Sonnet | Each supported Region: 200 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Sonnet. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude 3.5 Sonnet | ap-southeast-1: 40<br><br>Each of the other supported Regions: 100 | No | The maximum number of times that you can call model inference in one minute for Anthropic Claude 3.5 Sonnet. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference requests per minute for Anthropic Claude 3.7 Sonnet V1 | us-east-1: 250<br><br>us-east-2: 250<br><br>us-west-2: 250<br><br>eu-central-1: 100<br><br>eu-north-1: 100<br><br>eu-west-1: 100<br><br>eu-west-3: 100<br><br>Each of the other supported Regions: 50 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude 3.7 Sonnet V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude Haiku 4.5 | Each supported Region: 10,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Haiku 4.5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude Opus 4 V1 | Each supported Region: 200 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Opus 4 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude Opus 4.1 | Each supported Region: 50 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Opus 4.1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude Opus 4.5 | Each supported Region: 10,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Opus 4.5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude Opus 4.6 V1 | Each supported Region: 10,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Opus 4.6 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude Sonnet 4 V1 | Each supported Region: 200 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude Sonnet 4 V1 1M Context Length | Each supported Region: 5 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4 V1 1M Context Length. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude Sonnet 4.5 V1 | Each supported Region: 10,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length | Each supported Region: 1,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Anthropic Claude Sonnet 4.6 | Each supported Region: 10,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.6. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Cohere Embed V4 | Each supported Region: 2,000 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Cohere Embed V4. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for DeepSeek R1 V1 | Each supported Region: 200 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for DeepSeek R1 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Meta Llama 3.1 405B Instruct | Each supported Region: 400 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 3.1 405B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Meta Llama 3.1 70B Instruct | Each supported Region: 800 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 3.1 70B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Meta Llama 3.1 8B Instruct | Each supported Region: 1,600 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 3.1 8B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Meta Llama 3.2 1B Instruct | Each supported Region: 1,600 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 3.2 1B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Meta Llama 3.2 3B Instruct | Each supported Region: 1,600 | No | The maximum number of times that you can call model inference in one minute for Meta Llama 3.2 3B Instruct. The quota considers the combined sum of requests for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |
| Cross-region model inference requests per minute for Meta Llama 3.3 70B Instruct | Each supported Region: 800 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 3.3 70B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Meta Llama 4 Maverick V1 | Each supported Region: 800 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 4 Maverick V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Meta Llama 4 Scout V1 | Each supported Region: 800 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Meta Llama 4 Scout V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Mistral Pixtral Large 25.02 V1 | Each supported Region: 10 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Mistral Pixtral Large 25.02 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Conservative Upscale | Each supported Region: 4 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Conservative Upscale. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Stable Image Control Sketch | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Control Sketch. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Control Structure | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Control Structure. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Stable Image Creative Upscale | Each supported Region: 4 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Creative Upscale. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Erase Object | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Erase Object. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Stable Image Fast Upscale | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Fast Upscale. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Inpaint | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Inpaint. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Stable Image Outpaint | Each supported Region: 4 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Outpaint. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Remove Background | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Remove Background. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Stable Image Search and Recolor | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Search and Recolor. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Search and Replace | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Search and Replace. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Stable Image Style Guide | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Style Guide. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Stable Image Style Transfer | Each supported Region: 20 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Stable Image Style Transfer. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Twelve Labs Marengo | Each supported Region: 200 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Twelve Labs Marengo. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Twelve Labs Pegasus | Each supported Region: 120 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for Twelve Labs Pegasus. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for TwelveLabs Marengo Embed 3.0 | Each supported Region: 1,000 | Yes | The maximum number of cross-region requests that you can submit for model inference in one minute for TwelveLabs Marengo Embed 3.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference requests per minute for Writer AI Palmyra X4 V1 | Each supported Region: 10 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Writer AI Palmyra X4 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference requests per minute for Writer AI Palmyra X5 V1 | Each supported Region: 10 | No | The maximum number of cross-region requests that you can submit for model inference in one minute for Writer AI Palmyra X5 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Cross-region model inference tokens per minute for Amazon Nova 2 Lite | Each supported Region: 8,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova 2 Lite. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Amazon Nova 2 Omni | Each supported Region: 8,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova 2 Omni. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Amazon Nova 2 Pro Preview | Each supported Region: 1,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova 2 Pro Preview. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Amazon Nova Lite | us-east-1: 8,000,000<br><br>us-east-2: 8,000,000<br><br>us-west-1: 8,000,000<br><br>us-west-2: 8,000,000<br><br>Each of the other supported Regions: 400,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova Lite. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Amazon Nova Micro | us-east-1: 8,000,000<br><br>us-east-2: 8,000,000<br><br>us-west-2: 8,000,000<br><br>Each of the other supported Regions: 400,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova Micro. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Amazon Nova Premier V1 | Each supported Region: 2,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova Premier V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Amazon Nova Pro | Each supported Region: 2,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Amazon Nova Pro. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude 3 Haiku | us-east-1: 4,000,000<br><br>us-west-2: 4,000,000<br><br>ap-northeast-1: 400,000<br><br>ap-southeast-1: 400,000<br><br>Each of the other supported Regions: 600,000 | Yes | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Haiku. |
| Cross-region model inference tokens per minute for Anthropic Claude 3 Opus | Each supported Region: 800,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude 3 Opus. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude 3 Sonnet | Each supported Region: 400,000 | Yes | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Sonnet. |
| Cross-region model inference tokens per minute for Anthropic Claude 3.5 Sonnet | ap-southeast-1: 400,000<br><br>Each of the other supported Regions: 800,000 | Yes | The maximum number of tokens that you can submit for model inference in one minute for Anthropic Claude 3.5 Sonnet. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude 3.7 Sonnet V1 | us-east-1: 1,000,000<br><br>us-east-2: 1,000,000<br><br>us-west-2: 1,000,000<br><br>eu-central-1: 100,000<br><br>eu-north-1: 100,000<br><br>eu-west-1: 100,000<br><br>eu-west-3: 100,000<br><br>Each of the other supported Regions: 50,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude 3.7 Sonnet V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Fable 5 | Each supported Region: 4,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Fable 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Fable 5.1 | Each supported Region: 10,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Fable 5.1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Haiku 4.5 | Each supported Region: 5,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Haiku 4.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 4 V1 | Each supported Region: 200,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 4.1 | Each supported Region: 500,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 4.5 | Each supported Region: 2,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 4.6 V1 | Each supported Region: 3,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.6 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 4.7 | Each supported Region: 10,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.7. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 4.8 | Each supported Region: 30,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.8. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Opus 5 | Each supported Region: 30,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Sonnet 4 V1 | Each supported Region: 200,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Sonnet 4 V1 1M Context Length | Each supported Region: 1,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4 V1 1M Context Length. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Sonnet 4.5 V1 | Each supported Region: 5,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length | Each supported Region: 1,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Anthropic Claude Sonnet 4.6 | Each supported Region: 6,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Anthropic Claude Sonnet 5 | Each supported Region: 6,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Cohere Embed V4 | Each supported Region: 300,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Cohere Embed V4. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for DeepSeek R1 V1 | Each supported Region: 200,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for DeepSeek R1 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for GPT-5.6 Luna | Each supported Region: 20,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for GPT-5.6 Luna. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for GPT-5.6 Sol | Each supported Region: 10,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for GPT-5.6 Sol. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for GPT-5.6 Terra | Each supported Region: 20,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for GPT-5.6 Terra. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Grok 4.6 | Each supported Region: 10,000,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Grok 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Meta Llama 3.1 405B Instruct | Each supported Region: 800,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 3.1 405B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Meta Llama 3.1 70B Instruct | Each supported Region: 600,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 3.1 70B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Meta Llama 3.1 8B Instruct | Each supported Region: 600,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 3.1 8B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Meta Llama 3.2 1B Instruct | Each supported Region: 600,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 3.2 1B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Meta Llama 3.2 3B Instruct | Each supported Region: 600,000 | Yes | The maximum number of tokens that you can submit for model inference in one minute for Meta Llama 3.2 3B Instruct. The quota considers the combined sum of tokens for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |
| Cross-region model inference tokens per minute for Meta Llama 3.3 70B Instruct | Each supported Region: 600,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 3.3 70B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Meta Llama 4 Maverick V1 | Each supported Region: 600,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 4 Maverick V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Meta Llama 4 Scout V1 | Each supported Region: 600,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Meta Llama 4 Scout V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Mistral Pixtral Large 25.02 V1 | Each supported Region: 80,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Mistral Pixtral Large 25.02 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Cross-region model inference tokens per minute for Writer AI Palmyra X4 V1 | Each supported Region: 150,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Writer AI Palmyra X4 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cross-region model inference tokens per minute for Writer AI Palmyra X5 V1 | Each supported Region: 150,000 | Yes | The maximum number of cross-region tokens that you can submit for model inference in one minute for Writer AI Palmyra X5 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Custom models with a creating status per account | Each supported Region: 2 | Yes | The maximum number of custom models with a creating status. |
| DeleteAgent requests per second | Each supported Region: 2 | No | The maximum number of DeleteAgent API requests per second. |
| DeleteAgentActionGroup requests per second | Each supported Region: 2 | No | The maximum number of DeleteAgentActionGroup API requests per second. |
| DeleteAgentAlias requests per second | Each supported Region: 2 | No | The maximum number of DeleteAgentAlias API requests per second. |
| DeleteAgentVersion requests per second | Each supported Region: 2 | No | The maximum number of DeleteAgentVersion API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DisassociateAgentKnowledgeBase requests per second | Each supported Region: 4 | No | The maximum number of DisassociateAgentKnowledgeBase API requests per second. |
| Enabled action groups per agent | Each supported Region: 15 | Yes | The maximum number of action groups that you can enable in an Agent. |
| Endpoints per inference profile | Each supported Region: 5 | No | The maximum number of endpoints in an inference profile. An endpoint is defined by a model and the region that the invocation requests to the model are sent to. |
| GetAgent requests per second | Each supported Region: 15 | No | The maximum number of GetAgent API requests per second. |
| GetAgentActionGroup requests per second | Each supported Region: 20 | No | The maximum number of GetAgentActionGroup API requests per second. |
| GetAgentAlias requests per second | Each supported Region: 10 | No | The maximum number of GetAgentAlias API requests per second. |
| GetAgentKnowledgeBase requests per second | Each supported Region: 15 | No | The maximum number of GetAgentKnowledgeBase API requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetAgentVersion requests per second | Each supported Region: 10 | No | The maximum number of GetAgentVersion API requests per second. |
| Global cross-region model inference requests per minute for Amazon Nova 2 Lite | Each supported Region: 2,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Amazon Nova 2 Lite. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Global cross-region model inference requests per minute for Amazon Nova 2 Omni | Each supported Region: 2,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Amazon Nova 2 Omni. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference requests per minute for Amazon Nova 2 Pro Preview | Each supported Region: 100 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Amazon Nova 2 Pro Preview. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Global cross-region model inference requests per minute for Anthropic Claude Haiku 4.5 | Each supported Region: 10,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Haiku 4.5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference requests per minute for Anthropic Claude Opus 4.5 | Each supported Region: 10,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Opus 4.5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Global cross-region model inference requests per minute for Anthropic Claude Opus 4.6 V1 | Each supported Region: 10,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Opus 4.6 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference requests per minute for Anthropic Claude Sonnet 4 V1 | Each supported Region: 200 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Global cross-region model inference requests per minute for Anthropic Claude Sonnet 4.5 V1 | Each supported Region: 10,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference requests per minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length | Each supported Region: 1,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Global cross-region model inference requests per minute for Anthropic Claude Sonnet 4.6 | Each supported Region: 10,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.6. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference requests per minute for Cohere Embed V4 | Each supported Region: 2,000 | Yes | The maximum number of global cross-region requests that you can submit for model inference in one minute for Cohere Embed V4. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| Global cross-region model inference tokens per day for Amazon Nova 2 Lite | Each supported Region: 11,520,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Amazon Nova 2 Lite. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Amazon Nova 2 Omni | Each supported Region: 11,520,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Amazon Nova 2 Omni. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Amazon Nova 2 Pro Preview | Each supported Region: 1,440,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Amazon Nova 2 Pro Preview. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Fable 5 | Each supported Region: 5,760,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Fable 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Anthropic Claude Fable 5.1 | Each supported Region: 14,400,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Fable 5.1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Haiku 4.5 | Each supported Region: 7,200,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Haiku 4.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Anthropic Claude Opus 4.5 | Each supported Region: 2,880,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Opus 4.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Opus 4.6 V1 | Each supported Region: 4,320,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Opus 4.6 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Anthropic Claude Opus 4.7 | Each supported Region: 14,400,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Opus 4.7. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Opus 4.8 | Each supported Region: 43,200,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Opus 4.8. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Anthropic Claude Opus 5 | Each supported Region: 43,200,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Opus 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Sonnet 4 V1 | Each supported Region: 288,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Sonnet 4 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Anthropic Claude Sonnet 4.5 V1 | Each supported Region: 7,200,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Sonnet 4.5 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Sonnet 4.5 V1 1M Context Length | Each supported Region: 1,440,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Sonnet 4.5 V1 1M Context Length. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Anthropic Claude Sonnet 4.6 | Each supported Region: 8,640,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Sonnet 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for Anthropic Claude Sonnet 5 | Each supported Region: 8,640,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Anthropic Claude Sonnet 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Cohere Embed V4 | Each supported Region: 432,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Cohere Embed V4. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for GPT-5.6 Luna | Each supported Region: 28,800,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for GPT-5.6 Luna. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for GPT-5.6 Sol | Each supported Region: 14,400,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for GPT-5.6 Sol. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per day for GPT-5.6 Terra | Each supported Region: 28,800,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for GPT-5.6 Terra. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per day for Grok 4.6 | Each supported Region: 14,400,000,000 | No | The maximum number of global cross-region tokens that you can submit for model inference in one day for Grok 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Amazon Nova 2 Lite | Each supported Region: 8,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Amazon Nova 2 Lite. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Amazon Nova 2 Omni | Each supported Region: 8,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Amazon Nova 2 Omni. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Amazon Nova 2 Pro Preview | Each supported Region: 1,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Amazon Nova 2 Pro Preview. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Fable 5 | Each supported Region: 4,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Fable 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Anthropic Claude Fable 5.1 | Each supported Region: 10,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Fable 5.1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Haiku 4.5 | Each supported Region: 5,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Haiku 4.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Anthropic Claude Opus 4.5 | Each supported Region: 2,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Opus 4.6 V1 | Each supported Region: 3,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.6 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Anthropic Claude Opus 4.7 | Each supported Region: 10,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.7. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Opus 4.8 | Each supported Region: 30,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.8. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Anthropic Claude Opus 5 | Each supported Region: 30,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Opus 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Sonnet 4 V1 | Each supported Region: 200,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Anthropic Claude Sonnet 4.5 V1 | Each supported Region: 5,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length | Each supported Region: 1,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.5 V1 1M Context Length. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Anthropic Claude Sonnet 4.6 | Each supported Region: 6,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for Anthropic Claude Sonnet 5 | Each supported Region: 6,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Cohere Embed V4 | Each supported Region: 300,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Cohere Embed V4. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for GPT-5.6 Luna | Each supported Region: 20,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for GPT-5.6 Luna. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for GPT-5.6 Sol | Each supported Region: 10,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for GPT-5.6 Sol. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Global cross-region model inference tokens per minute for GPT-5.6 Terra | Each supported Region: 20,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for GPT-5.6 Terra. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global cross-region model inference tokens per minute for Grok 4.6 | Each supported Region: 10,000,000 | Yes | The maximum number of global cross-region tokens that you can submit for model inference in one minute for Grok 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Imported models per account | Each supported Region: 3 | Yes | The maximum number of imported models in an account. |
| Inference profiles per account | Each supported Region: 1,000 | Yes | The maximum number of inference profiles in an account. |
| ListAgentActionGroups requests per second | Each supported Region: 10 | No | The maximum number of ListAgentActionGroups API requests per second. |
| ListAgentAliases requests per second | Each supported Region: 10 | No | The maximum number of ListAgentAliases API requests per second. |
| ListAgentKnowledgeBases requests per second | Each supported Region: 10 | No | The maximum number of ListAgentKnowledgeBases API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ListAgentVersions requests per second | Each supported Region: 10 | No | The maximum number of ListAgentVersions API requests per second. |
| ListAgents requests per second | Each supported Region: 10 | No | The maximum number of ListAgents API requests per second. |
| Minimum number of records per batch inference job for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Amazon Nova 2 Multimodal Embeddings V1. |
| Minimum number of records per batch inference job for Amazon Nova Premier | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Amazon Nova Premier. |
| Minimum number of records per batch inference job for Claude 3 Haiku | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude 3 Haiku. |
| Minimum number of records per batch inference job for Claude 3 Opus | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude 3 Opus. |
| Minimum number of records per batch inference job for Claude 3 Sonnet | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude 3 Sonnet. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for Claude 3.5 Haiku | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude 3.5 Haiku. |
| Minimum number of records per batch inference job for Claude 3.5 Sonnet | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude 3.5 Sonnet. |
| Minimum number of records per batch inference job for Claude 3.5 Sonnet v2 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude 3.5 Sonnet v2. |
| Minimum number of records per batch inference job for Claude 3.7 Sonnet | Each supported Region: 100 | Yes | The minimum number of records across all input files in a batch inference job for Claude 3.7 Sonnet. |
| Minimum number of records per batch inference job for Claude Haiku 4.5 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude Haiku 4.5. |
| Minimum number of records per batch inference job for Claude Opus 4.5 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude Opus 4.5. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for Claude Opus 4.6 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude Opus 4.6. |
| Minimum number of records per batch inference job for Claude Opus 5 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude Opus 5. |
| Minimum number of records per batch inference job for Claude Sonnet 4 | Each supported Region: 100 | Yes | The minimum number of records across all input files in a batch inference job for Claude Sonnet 4. |
| Minimum number of records per batch inference job for Claude Sonnet 4.5. | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude Sonnet 4.5. |
| Minimum number of records per batch inference job for Claude Sonnet 4.6 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Claude Sonnet 4.6. |
| Minimum number of records per batch inference job for DeepSeek V3.2 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for DeepSeek V3.2. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for DeepSeek v3 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for DeepSeek v3. |
| Minimum number of records per batch inference job for Devstral 2 123B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Devstral 2 123B. |
| Minimum number of records per batch inference job for GLM 4.7 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for GLM 4.7. |
| Minimum number of records per batch inference job for GLM 4.7 Flash | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for GLM 4.7 Flash. |
| Minimum number of records per batch inference job for GLM 5 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for GLM 5. |
| Minimum number of records per batch inference job for Gemma 3 12B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Gemma 3 12B. |
| Minimum number of records per batch inference job for Gemma 3 27B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Gemma 3 27B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for Gemma 3 4B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Gemma 3 4B. |
| Minimum number of records per batch inference job for Kimi K2 Thinking | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Kimi K2 Thinking. |
| Minimum number of records per batch inference job for Kimi K2.5 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Kimi K2.5. |
| Minimum number of records per batch inference job for Llama 3.1 405B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.1 405B Instruct. |
| Minimum number of records per batch inference job for Llama 3.1 70B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.1 70B Instruct. |
| Minimum number of records per batch inference job for Llama 3.1 8B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.1 8B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for Llama 3.2 11B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.2 11B Instruct. |
| Minimum number of records per batch inference job for Llama 3.2 1B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job Llama 3.2 1B Instruct. |
| Minimum number of records per batch inference job for Llama 3.2 3B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.2 3B Instruct. |
| Minimum number of records per batch inference job for Llama 3.2 90B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.2 90B Instruct. |
| Minimum number of records per batch inference job for Llama 3.3 70B Instruct | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 3.3 70B Instruct. |
| Minimum number of records per batch inference job for Llama 4 Maverick | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 4 Maverick. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for Llama 4 Scout | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Llama 4 Scout. |
| Minimum number of records per batch inference job for Magistral Small 2509 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Magistral Small 2509. |
| Minimum number of records per batch inference job for MiniMax M2 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for MiniMax M2. |
| Minimum number of records per batch inference job for MiniMax M2.1 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for MiniMax M2.1. |
| Minimum number of records per batch inference job for MiniMax M2.5 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for MiniMax M2.5. |
| Minimum number of records per batch inference job for Ministral 3 14B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Ministral 3 14B. |
| Minimum number of records per batch inference job for Ministral 3 8B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Ministral 3 8B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for Ministral 3B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Ministral 3B. |
| Minimum number of records per batch inference job for Mistral Large 2 (24.07) | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Mistral Large 2 (24.07). |
| Minimum number of records per batch inference job for Mistral Large 3 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Mistral Large 3. |
| Minimum number of records per batch inference job for Mistral Small | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Mistral Small. |
| Minimum number of records per batch inference job for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for NVIDIA Nemotron 3 Super 120B A12B. |
| Minimum number of records per batch inference job for NVIDIA Nemotron Nano 12B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for NVIDIA Nemotron Nano 12B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for NVIDIA Nemotron Nano 3 30B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for NVIDIA Nemotron Nano 3 30B. |
| Minimum number of records per batch inference job for NVIDIA Nemotron Nano 9B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for NVIDIA Nemotron Nano 9B. |
| Minimum number of records per batch inference job for Nova 2 Lite | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Nova 2 Lite. |
| Minimum number of records per batch inference job for Nova Lite V1 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Nova Lite V1. |
| Minimum number of records per batch inference job for Nova Micro V1 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Nova Micro V1. |
| Minimum number of records per batch inference job for Nova Pro V1 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Nova Pro V1. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of records per batch inference job for OpenAI GPT OSS 120b | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for OpenAI GPT OSS 120b. |
| Minimum number of records per batch inference job for OpenAI GPT OSS 20b | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for OpenAI GPT OSS 20b. |
| Minimum number of records per batch inference job for OpenAI GPT OSS Safeguard 120b | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for OpenAI GPT OSS Safeguard 120b. |
| Minimum number of records per batch inference job for OpenAI GPT OSS Safeguard 20b | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for OpenAI GPT OSS Safeguard 20b. |
| Minimum number of records per batch inference job for Qwen3 235B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 235B. |
| Minimum number of records per batch inference job for Qwen3 32B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 32B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Minimum number of records per batch inference job for Qwen3 Coder 30B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 Coder 30B. |
| Minimum number of records per batch inference job for Qwen3 Coder 480B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 Coder 480B. |
| Minimum number of records per batch inference job for Qwen3 Coder Next | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 Coder Next. |
| Minimum number of records per batch inference job for Qwen3 Next 80B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 Next 80B. |
| Minimum number of records per batch inference job for Qwen3 VL 235B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Qwen3 VL 235B. |
| Minimum number of records per batch inference job for Titan Multimodal Embeddings G1 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Titan Multimodal Embeddings G1. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Minimum number of records per batch inference job for Titan Text Embeddings V2 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Titan Text Embeddings V2. |
| Minimum number of records per batch inference job for Voxtral Mini 3B 2507 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Voxtral Mini 3B 2507. |
| Minimum number of records per batch inference job for Voxtral Small 24B 2507 | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Voxtral Small 24B 2507. |
| Minimum number of records per batch inference job for Writer Palmyra Vision 7B | Each supported Region: 100 | No | The minimum number of records across all input files in a batch inference job for Writer Palmyra Vision 7B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max latency-optimized tokens per day for Amazon Nova Pro V1 | Each supported Region: 57,600,000 | No | The maximum number of tokens that you can submit for latency-optimized model inference in one day for Amazon Nova Pro V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Model invocation max tokens per day for AI21 Labs Jamba 1.5 Large (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for AI21 Labs Jamba 1.5 Large. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for AI21 Labs Jamba 1.5 Mini (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for AI21 Labs Jamba 1.5 Mini. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Amazon Nova 2 Lite (doubled for cross-region calls) | Each supported Region: 5,760,000,000 | No | Daily maximum tokens for model inference for Amazon Nova 2 Lite. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Amazon Nova 2 Omni (doubled for cross-region calls) | Each supported Region: 5,760,000,000 | No | Daily maximum tokens for model inference for Amazon Nova 2 Omni. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Amazon Nova 2 Pro Preview (doubled for cross-region calls) | Each supported Region: 720,000,000 | No | Daily maximum tokens for model inference for Amazon Nova 2 Pro Preview. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Amazon Nova Lite (doubled for cross-region calls) | Each supported Region: 5,760,000,000 | No | Daily maximum tokens for model inference for Amazon Nova Lite. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Amazon Nova Micro (doubled for cross-region calls) | Each supported Region: 5,760,000,000 | No | Daily maximum tokens for model inference for Amazon Nova Micro. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Amazon Nova Premier V1 (doubled for cross-region calls) | Each supported Region: 1,440,000,000 | No | Daily maximum tokens for model inference for Amazon Nova Premier V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Amazon Nova Pro (doubled for cross-region calls) | Each supported Region: 1,440,000,000 | No | Daily maximum tokens for model inference for Amazon Nova Pro. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude 3 Haiku (doubled for cross-region calls) | us-east-1: 2,880,000,000<br><br>us-west-2: 2,880,000,000<br><br>ap-northeast-1: 288,000,000<br><br>ap-southeast-1: 288,000,000<br><br>Each of the other supported Regions: 432,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude 3 Haiku. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude 3.5 Haiku (doubled for cross-region calls) | Each supported Region: 2,880,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude 3.5 Haiku. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude 3.5 Sonnet V1 (doubled for cross-region calls) | Each supported Region: 2,880,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude 3.5 Sonnet V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude 3.5 Sonnet V2 (doubled for cross-region calls) | Each supported Region: 576,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude 3.5 Sonnet V2. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude 3.7 Sonnet V1 (doubled for cross-region calls) | Each supported Region: 720,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude 3.7 Sonnet V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Fable 5 (doubled for cross-region calls) | Each supported Region: 2,880,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Fable 5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Fable 5.1 (doubled for cross-region calls) | Each supported Region: 7,200,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Fable 5.1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Haiku 4.5 (doubled for cross-region calls) | Each supported Region: 3,600,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Haiku 4.5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Opus 4 V1 (doubled for cross-region calls) | Each supported Region: 144,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 4 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Opus 4.1 (doubled for cross-region calls) | Each supported Region: 360,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 4.1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Opus 4.5 (doubled for cross-region calls) | Each supported Region: 1,440,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 4.5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Opus 4.6 V1 (doubled for cross-region calls) | Each supported Region: 2,160,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 4.6 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Opus 4.7 (doubled for cross-region calls) | Each supported Region: 7,200,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 4.7. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Opus 4.8 (doubled for cross-region calls) | Each supported Region: 21,600,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 4.8. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Opus 5 (doubled for cross-region calls) | Each supported Region: 21,600,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Opus 5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Sonnet 4 V1 (doubled for cross-region calls) | Each supported Region: 144,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Sonnet 4 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Sonnet 4 V1 1M Context Length (doubled for cross-region calls) | Each supported Region: 720,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Sonnet 4 V1 1M Context Length. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Sonnet 4.5 V1 (doubled for cross-region calls) | Each supported Region: 3,600,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Sonnet 4.5 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Sonnet 4.5 V1 1M Context Length (doubled for cross-region calls) | Each supported Region: 720,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Sonnet 4.5 V1 1M Context Length. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Anthropic Claude Sonnet 4.6 (doubled for cross-region calls) | Each supported Region: 4,320,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Sonnet 4.6. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Anthropic Claude Sonnet 5 (doubled for cross-region calls) | Each supported Region: 4,320,000,000 | No | Daily maximum tokens for model inference for Anthropic Claude Sonnet 5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Cohere Embed V4 (doubled for cross-region calls) | Each supported Region: 216,000,000 | No | Daily maximum tokens for model inference for Cohere Embed V4. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for DeepSeek R1 V1 (doubled for cross-region calls) | Each supported Region: 144,000,000 | No | Daily maximum tokens for model inference for DeepSeek R1 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for DeepSeek V3 V1 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for DeepSeek V3 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for DeepSeek V3.2 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for DeepSeek V3.2. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for GPT OSS Safeguard 120B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for GPT OSS Safeguard 120B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for GPT OSS Safeguard 20B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for GPT OSS Safeguard 20B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for GPT-5.6 Luna (doubled for cross-region calls) | Each supported Region: 14,400,000,000 | No | Daily maximum tokens for model inference for GPT-5.6 Luna. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for GPT-5.6 Sol (doubled for cross-region calls) | Each supported Region: 7,200,000,000 | No | Daily maximum tokens for model inference for GPT-5.6 Sol. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for GPT-5.6 Terra (doubled for cross-region calls) | Each supported Region: 14,400,000,000 | No | Daily maximum tokens for model inference for GPT-5.6 Terra. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Gemma 3 12B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Gemma 3 12B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Gemma 3 27B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Gemma 3 27B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Gemma 3 4B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Gemma 3 4B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Grok 4.6 (doubled for cross-region calls) | Each supported Region: 7,200,000,000 | No | Daily maximum tokens for model inference for Grok 4.6. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Kimi K2 Thinking (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Kimi K2 Thinking. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Magistral Small 1.2 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Magistral Small 1.2. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Meta Llama 3.2 11B Instruct (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Meta Llama 3.2 11B Instruct. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Meta Llama 3.2 1B Instruct (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Meta Llama 3.2 1B Instruct. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Meta Llama 3.2 3B Instruct (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Meta Llama 3.2 3B Instruct. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Meta Llama 3.2 90B Instruct (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Meta Llama 3.2 90B Instruct. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Meta Llama 4 Maverick V1 (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Meta Llama 4 Maverick V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Meta Llama 4 Scout V1 (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Meta Llama 4 Scout V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for MiniMax M2.5 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for MiniMax M2.5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Minimax M2 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Minimax M2. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Minimax M2.1 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Minimax M2.1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Ministral 14B 3.0 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Ministral 14B 3.0. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Ministral 3B 3.0 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Ministral 3B 3.0. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Ministral 8B 3.0 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Ministral 8B 3.0. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Mistral AI Mistral 7B Instruct (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Mistral AI Mistral 7B Instruct. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Mistral AI Mistral Large (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Mistral AI Mistral Large. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Mistral AI Mistral Small (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Mistral AI Mistral Small. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Mistral AI Mixtral 8X7B Instruct (doubled for cross-region calls) | Each supported Region: 432,000,000 | No | Daily maximum tokens for model inference for Mistral AI Mixtral 8X7B Instruct. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Mistral Devstral 2 123b (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Mistral Devstral 2 123b. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Mistral Large 3 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Mistral Large 3. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Mistral Pixtral Large 25.02 V1 (doubled for cross-region calls) | Each supported Region: 57,600,000 | No | Daily maximum tokens for model inference for Mistral Pixtral Large 25.02 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Moonshot AI Kimi K2.5 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Moonshot AI Kimi K2.5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for NVIDIA Nemotron 3 Super 120B A12B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for NVIDIA Nemotron 3 Super 120B A12B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for NVIDIA Nemotron Nano 2 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for NVIDIA Nemotron Nano 2. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for NVIDIA Nemotron Nano 2 VL (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for NVIDIA Nemotron Nano 2 VL. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Nemotron Nano 3 30B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Nemotron Nano 3 30B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for OpenAI GPT OSS 120B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for OpenAI GPT OSS 120B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for OpenAI GPT OSS 20B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for OpenAI GPT OSS 20B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Qwen3 235B a22b 2507 V1 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 235B a22b 2507 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Qwen3 32B V1 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 32B V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Qwen3 Coder 30B a3b V1 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 Coder 30B a3b V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Qwen3 Coder 480B a35b V1 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 Coder 480B a35b V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Qwen3 Coder Next (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 Coder Next. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Qwen3 Next 80B A3B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 Next 80B A3B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Qwen3 VL 235B A22B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Qwen3 VL 235B A22B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Voxtral Mini 1.0 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Voxtral Mini 1.0. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Voxtral Small 1.0 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Voxtral Small 1.0. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Writer AI Palmyra X4 V1 (doubled for cross-region calls) | Each supported Region: 108,000,000 | No | Daily maximum tokens for model inference for Writer AI Palmyra X4 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Writer AI Palmyra X5 V1 (doubled for cross-region calls) | Each supported Region: 108,000,000 | No | Daily maximum tokens for model inference for Writer AI Palmyra X5 V1. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Writer Palmyra Vision 7B (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Writer Palmyra Vision 7B. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model invocation max tokens per day for Z.ai GLM 5 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Z.ai GLM 5. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model invocation max tokens per day for Z.ai GLM-4.7 (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Z.ai GLM-4.7. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Model invocation max tokens per day for Z.ai GLM-4.7 Flash (doubled for cross-region calls) | Each supported Region: 144,000,000,000 | No | Daily maximum tokens for model inference for Z.ai GLM-4.7 Flash. Combines sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. Doubled for cross-region calls; not applicable in case of approved TPM increase. |
| Model units no-commitment Provisioned Throughputs across base models | Each supported Region: 0 | Yes | The maximum number of model units that can be distributed across no-commitment Provisioned Throughputs for base models |
| Model units no-commitment Provisioned Throughputs across custom models | Each supported Region: 0 | Yes | The maximum number of model units that can be distributed across no-commitment Provisioned Throughputs for custom models |
| Model units per provisioned model for AI21 Labs Jurassic-2 Mid | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for AI21 Labs Jurassic-2 Mid. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for AI21 Labs Jurassic-2 Ultra | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for AI21 Labs Jurassic-2 Ultra. |
| Model units per provisioned model for Amazon Nova 2 Lite V1.0 256K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Nova 2 Lite V1.0 256K. |
| Model units per provisioned model for Amazon Nova Canvas | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Nova Canvas. |
| Model units per provisioned model for Amazon Titan Embeddings G1 - Text | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Embeddings G1 - Text. |
| Model units per provisioned model for Amazon Titan Image Generator G1 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Image Generator G1. |
| Model units per provisioned model for Amazon Titan Image Generator G2 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Image Generator G2. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for Amazon Titan Lite V1 4K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Text Lite V1 4K. |
| Model units per provisioned model for Amazon Titan Multimodal Embeddings G1 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Multimodal Embeddings G1. |
| Model units per provisioned model for Amazon Titan Text Embeddings V2 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Text Embeddings V2. |
| Model units per provisioned model for Amazon Titan Text G1 - Express 8K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Text G1 - Express 8K. |
| Model units per provisioned model for Amazon Titan Text Premier V1 32K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Amazon Titan Text Premier V1 32K. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for Anthropic Claude 3 Haiku 200K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3 Haiku 200K. |
| Model units per provisioned model for Anthropic Claude 3 Haiku 48K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3 Haiku 48K. |
| Model units per provisioned model for Anthropic Claude 3 Sonnet 200K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3 Sonnet 200K. |
| Model units per provisioned model for Anthropic Claude 3 Sonnet 28K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3 Sonnet 28K. |
| Model units per provisioned model for Anthropic Claude 3.5 Haiku 16K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.5 Haiku 16K. |
| Model units per provisioned model for Anthropic Claude 3.5 Haiku 200K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.5 Haiku 200K. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for Anthropic Claude 3.5 Haiku 64K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.5 Haiku 64K. |
| Model units per provisioned model for Anthropic Claude 3.5 Sonnet 18K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.5 Sonnet 18K. |
| Model units per provisioned model for Anthropic Claude 3.5 Sonnet 200K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.5 Sonnet 200K. |
| Model units per provisioned model for Anthropic Claude 3.5 Sonnet 51K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.5 Sonnet 51K. |
| Model units per provisioned model for Anthropic Claude 3.7 V1.0 Sonnet 18K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.7 V1.0 Sonnet 18K. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for Anthropic Claude 3.7 V1.0 Sonnet 200K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.7 V1.0 Sonnet 200K. |
| Model units per provisioned model for Anthropic Claude 3.7 V1.0 Sonnet 51K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude 3.7 V1.0 Sonnet 51K. |
| Model units per provisioned model for Anthropic Claude Instant V1 100K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude Instant V1 100K. |
| Model units per provisioned model for Anthropic Claude V2 100K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude V2 100K. |
| Model units per provisioned model for Anthropic Claude V2 18K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude V2 18K. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Model units per provisioned model for Anthropic Claude V2.1 18K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude V2.1 18K. |
| Model units per provisioned model for Anthropic Claude V2.1 200K | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Anthropic Claude V2.1 200K. |
| Model units per provisioned model for Cohere Command | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Cohere Command. |
| Model units per provisioned model for Cohere Command Light | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Cohere Command Light. |
| Model units per provisioned model for Cohere Command R | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Cohere Command R 128k. |
| Model units per provisioned model for Cohere Command R Plus | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Cohere Command R Plus 128k. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for Cohere Embed English | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Cohere Embed English. |
| Model units per provisioned model for Cohere Embed Multilingual | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Cohere Embed Multilingual. |
| Model units per provisioned model for Meta Llama 2 13B | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 2 13B. |
| Model units per provisioned model for Meta Llama 2 70B | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 2 70B. |
| Model units per provisioned model for Meta Llama 2 Chat 13B | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 2 Chat 13B. |
| Model units per provisioned model for Meta Llama 2 Chat 70B | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 2 Chat 70B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for Meta Llama 3 70B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3 70B Instruct. |
| Model units per provisioned model for Meta Llama 3 8B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3 8B Instruct. |
| Model units per provisioned model for Meta Llama 3.1 70B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3.1 70B Instruct. |
| Model units per provisioned model for Meta Llama 3.1 8B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3.1 8B Instruct. |
| Model units per provisioned model for Meta Llama 3.2 11B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3.2 11B Instruct. |
| Model units per provisioned model for Meta Llama 3.2 1B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3.2 1B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Model units per provisioned model for Meta Llama 3.2 3B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3.2 3B Instruct. |
| Model units per provisioned model for Meta Llama 3.2 90B Instruct | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Meta Llama 3.2 90B Instruct. |
| Model units per provisioned model for Mistral Large 2407 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Mistral Large 2407. |
| Model units per provisioned model for Mistral Small | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Mistral Small. |
| Model units per provisioned model for Stability.ai Stable Diffusion XL 0.8 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Stability.ai Stable Diffusion XL 0.8 |
| Model units per provisioned model for Stability.ai Stable Diffusion XL 1.0 | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for Stability.ai Stable Diffusion XL 1.0. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for the 128k context length variant for Amazon Nova Micro | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for the 128k context length variant for Amazon Nova Micro |
| Model units per provisioned model for the 24k context length variant for Amazon Nova Lite | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for the 24k context length variant for Amazon Nova Lite |
| Model units per provisioned model for the 24k context length variant for Amazon Nova Micro | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for the 24k context length variant for Amazon Nova Micro |
| Model units per provisioned model for the 24k context length variant for Amazon Nova Pro | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for the 24k context length variant for Amazon Nova Pro |
| Model units per provisioned model for the 300k context length variant for Amazon Nova Lite | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for the 300k context length variant for Amazon Nova Lite |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units per provisioned model for the 300k context length variant for Amazon Nova Pro | Each supported Region: 0 | Yes | The maximum number of model units that can be allotted to a provisioned model for the 300k context length variant for Amazon Nova Pro. |
| Model units, with commitment, for Provisioned Throughout created for Meta Llama 4 Scout 17B Instruct 10M | Each supported Region: 0 | Yes | The maximum number of model units that can be allocated to a Provisioned Throughput created for Meta Llama 4 Scout 17B Instruct 10M, with commitment. |
| Model units, with commitment, for Provisioned Throughout created for Meta Llama 4 Scout 17B Instruct 128K | Each supported Region: 0 | Yes | The maximum number of model units that can be allocated to a Provisioned Throughput created for Meta Llama 4 Scout 17B Instruct 128K, with commitment. |
| Model units, with commitment, for Provisioned Throughout created for Meta Maverick 4 Scout 17B Instruct 128K | Each supported Region: 0 | Yes | The maximum number of model units that can be allocated to a Provisioned Throughput created for Meta Llama 4 Maverick 17B Instruct 128K, with commitment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Model units, with commitment, for Provisioned Throughout created for Meta Maverick 4 Scout 17B Instruct 1M | Each supported Region: 0 | Yes | The maximum number of model units that can be allocated to a Provisioned Throughput created for Meta Llama 4 Maverick 17B Instruct 1M, with commitment. |
| No-commitment model units for Provisioned Throughput created for base model Amazon Nova 2 Lite V1.0 256K | Each supported Region: 0 | No | The maximum number of model units that can be allocated to a Provisioned Throughput created for the base model Amazon Nova 2 Lite V1.0 256K, with no commitment |
| No-commitment model units for Provisioned Throughput created for custom model Amazon Nova 2 Lite V1.0 256K | Each supported Region: 0 | No | The maximum number of model units that can be allocated to a Provisioned Throughput created for the custom model Amazon Nova 2 Lite V1.0 256K, with no commitment. |
| Number of custom prompt routers per account | Each supported Region: 500 | No | The maximum number of custom prompt routers that you can create per account per region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-Demand, latency-optimized model inference requests per minute for Amazon Nova Pro V1 | Each supported Region: 10 | No | The maximum number of on-demand, latency-optimized requests that you can submit for model inference in one minute for Amazon Nova Pro V1. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-Demand, latency-optimized model inference requests per minute for Meta Llama 3.1 405B Instruct | Each supported Region: 100 | No | The maximum number of on-demand, latency-optimized requests that you can submit for model inference in one minute for Meta Llama 3.1 405B Instruct. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| On-Demand, latency-optimized model inference requests per minute for Meta Llama 3.1 70B Instruct | Each supported Region: 100 | No | The maximum number of on-demand, latency-optimized requests that you can submit for model inference in one minute for Meta Llama 3.1 70B Instruct. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-Demand, latency-optimized model inference tokens per minute for Amazon Nova Pro V1 | Each supported Region: 40,000 | No | The maximum number of on-demand, latency-optimized tokens that you can submit for model inference in one minute for Amazon Nova Pro V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| On-Demand, latency-optimized model inference tokens per minute for Meta Llama 3.1 405B Instruct | Each supported Region: 40,000 | No | The maximum number of on-demand, latency-optimized tokens that you can submit for model inference in one minute for Meta Llama 3.1 405B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-Demand, latency-optimized model inference tokens per minute for Meta Llama 3.1 70B Instruct | Each supported Region: 40,000 | No | The maximum number of on-demand, latency-optimized tokens that you can submit for model inference in one minute for Meta Llama 3.1 70B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand InvokeModel async concurrent requests for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 30 | No | The maximum number of async concurrent model inference requests that you can submit for Amazon Nova 2 Multimodal Embeddings V1. The quota considers the combined sum of Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream. |
| On-demand InvokeModel async concurrent requests for TwelveLabs Marengo Embed 3.0 | Each supported Region: 10 | Yes | The maximum number of async concurrent model inference requests that you can submit for TwelveLabs Marengo Embed 3.0. The quota considers the combined sum of Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand InvokeModel concurrent requests for Amazon Nova 2 Sonic | Each supported Region: 20 | No | The maximum number of concurrent model inference requests that you can submit for Amazon Nova 2 Sonic. The quota considers the combined sum of Converse, ConverseStream, InvokeModel, and InvokeModelWithResponseStream. |
| On-demand InvokeModel concurrent requests for Amazon Nova Reel1.0 | Each supported Region: 10 | No | The maximum number of concurrent model inference requests that you can submit for Amazon Nova Reel 1.0. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand InvokeModel concurrent requests for Amazon Nova Reel1.1 | Each supported Region: 3 | No | The maximum number of concurrent model inference requests that you can submit for Amazon Nova Reel 1.1. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand InvokeModel concurrent requests for Amazon Nova Sonic | Each supported Region: 20 | No | The maximum number of concurrent requests that you can submit for model inference for Amazon Nova Sonic. |
| On-demand InvokeModel concurrent requests for Twelve Labs Marengo | Each supported Region: 30 | No | The maximum number of concurrent InvokeModel requests that you can submit for model inference for Twelve Labs Marengo. |
| On-demand InvokeModel concurrent requests for Twelve Labs Pegasus | Each supported Region: 30 | Yes | The maximum number of concurrent InvokeModel requests that you can submit for model inference for Twelve Labs Pegasus. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand latency-optimized model inference requests per minute for Anthropic Claude 3.5 Haiku | Each supported Region: 100 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3.5 Haiku, if latency optimization is configured. |
| On-demand latency-optimized model inference tokens per minute for Anthropic Claude 3.5 Haiku | Each supported Region: 500,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3.5 Haiku, if latency optimization is configured. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference concurrent requests for Luma Ray V2 | Each supported Region: 1 | No | The maximum number of concurrent requests that you can submit for model inference for Luma Ray V2. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference requests per minute for AI21 Labs Jamba 1.5 Large | Each supported Region: 100 | No | The maximum number of times that you can call model inference in one minute for AI21 Labs Jamba 1.5 Large. The quota considers the combined sum of requests for Converse and InvokeModel |
| On-demand model inference requests per minute for AI21 Labs Jamba 1.5 Mini | Each supported Region: 100 | No | The maximum number of times that you can call model inference in one minute for AI21 Labs Jamba 1.5 Mini. The quota considers the combined sum of requests for Converse and InvokeModel |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 2,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Amazon Nova 2 Multimodal Embeddings V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Amazon Nova Canvas | Each supported Region: 100 | No | The maximum number of requests that you can submit for model inference in one minute for Amazon Nova Canvas. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Amazon Nova Lite | us-east-1: 2,000<br>eu-west-2: 2,000<br>Each of the other supported Regions: 200 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Amazon Nova Lite. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Amazon Nova Micro | us-east-1: 2,000<br>eu-west-2: 2,000<br>Each of the other supported Regions: 200 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Amazon Nova Micro. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Amazon Nova Pro | Each supported Region: 250 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Amazon Nova Pro. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Amazon Rerank 1.0 | Each supported Region: 200 | No | The maximum number of times that you can call InvokeModel in one minute for Amazon Rerank 1.0. |
| On-demand model inference requests per minute for Amazon Titan Image Generator G1 | Each supported Region: 60 | No | The maximum number of times that you can call InvokeModel in one minute for Amazon Titan Image Generator G1. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Amazon Titan Image Generator G1 V2 | Each supported Region: 60 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Amazon Titan Image Generator G1 V2. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Amazon Titan Multimodal Embeddings G1 | Each supported Region: 2,000 | No | The maximum number of times that you can call InvokeModel in one minute for Amazon Titan Multimodal Embeddings G1. |
| On-demand model inference requests per minute for Amazon Titan Text Embeddings | Each supported Region: 2,000 | No | The maximum number of times that you can call InvokeModel in one minute for Amazon Titan Text Embeddings |
| On-demand model inference requests per minute for Amazon Titan Text Embeddings V2 | Each supported Region: 6,000 | No | The maximum number of times that you can call InvokeModel in one minute for Amazon Titan Text Embeddings V2 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Amazon Titan Text Express | Each supported Region: 400 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Amazon Titan Text Express |
| On-demand model inference requests per minute for Amazon Titan Text Premier | Each supported Region: 100 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Amazon Titan Text Premier |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Anthropic Claude 3 Haiku | us-east-1: 1,000<br><br>us-west-2: 1,000<br><br>ap-northeast-1: 200<br><br>ap-southeast-1: 200<br><br>Each of the other supported Regions: 400 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Haiku. |
| On-demand model inference requests per minute for Anthropic Claude 3 Opus | Each supported Region: 50 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Anthropic Claude 3 Opus. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Anthropic Claude 3 Sonnet | Each supported Region: 100 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Sonnet. |
| On-demand model inference requests per minute for Anthropic Claude 3.5 Haiku | us-west-1: 400<br><br>Each of the other supported Regions: 1,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Anthropic Claude 3.5 Haiku. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Anthropic Claude 3.5 Sonnet | ap-southeast-1: 20<br><br>Each of the other supported Regions: 50 | No | The maximum number of times that you can call model inference in one minute for Anthropic Claude 3.5 Sonnet. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference requests per minute for Anthropic Claude 3.5 Sonnet V2 | Each supported Region: 50 | No | The maximum number of times that you can call model inference in one minute for Anthropic Claude 3.5 Sonnet V2. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Anthropic Claude 3.7 Sonnet V1 | Each supported Region: 125 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Anthropic Claude 3.7 Sonnet V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Anthropic Claude Opus 4.6 V1 | Each supported Region: 5,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Anthropic Claude Opus 4.6 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Anthropic Claude Sonnet 4.6 | Each supported Region: 5,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.6. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Cohere Command R | Each supported Region: 400 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Cohere Command R 128k. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Cohere Command R Plus | Each supported Region: 400 | No | The maximum number of times that you can call model inference in one minute. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Cohere Command R Plus 128k. |
| On-demand model inference requests per minute for Cohere Embed English | Each supported Region: 2,000 | No | The maximum number of times that you can call InvokeModel in one minute for Cohere Embed English. |
| On-demand model inference requests per minute for Cohere Embed Multilingual | Each supported Region: 2,000 | No | The maximum number of times that you can call InvokeModel in one minute for Cohere Embed Multilingual. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Cohere Embed V4 | Each supported Region: 1,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Cohere Embed V4. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Cohere Rerank 3.5 | Each supported Region: 250 | No | The maximum number of times that you can call InvokeModel in one minute for Cohere Rerank 3.5. |
| On-demand model inference requests per minute for DeepSeek V3 V1 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for DeepSeek V3 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for DeepSeek V3.2 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for DeepSeek V3.2. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for GPT OSS Safeguard 120B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for GPT OSS Safeguard 120B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for GPT OSS Safeguard 20B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for GPT OSS Safeguard 20B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Gemma 3 12B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Gemma 3 12B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Gemma 3 27B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Gemma 3 27B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Gemma 3 4B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Gemma 3 4B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Kimi K2 Thinking | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Kimi K2 Thinking. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Magistral Small 1.2 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Magistral Small 1.2. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Meta Llama 3 70B Instruct | Each supported Region: 400 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3 70B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Meta Llama 3 8B Instruct | Each supported Region: 800 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3 8B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Meta Llama 3.1 405B Instruct | Each supported Region: 200 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3.1 405B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Meta Llama 3.1 70B Instruct | Each supported Region: 400 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3.1 70B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Meta Llama 3.1 8B Instruct | Each supported Region: 800 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3.1 8B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Meta Llama 3.2 11B Instruct | Each supported Region: 400 | No | The maximum number of times that you can call model inference in one minute for Meta Llama 3.2 11B Instruct. The quota considers the combined sum of requests for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Meta Llama 3.2 1B Instruct | Each supported Region: 800 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3.2 1B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Meta Llama 3.2 3B Instruct | Each supported Region: 800 | No | The maximum number of times that you can call model inference in one minute for Meta Llama 3.2 3B Instruct. The quota considers the combined sum of requests for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Meta Llama 3.2 90B Instruct | Each supported Region: 400 | No | The maximum number of times that you can call model inference in one minute for Meta Llama 3.2 90B Instruct. The quota considers the combined sum of requests for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |
| On-demand model inference requests per minute for Meta Llama 3.3 70B Instruct | Each supported Region: 400 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Meta Llama 3.3 70B Instruct. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for MiniMax M2.5 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for MiniMax M2.5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Minimax M2 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Minimax M2. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Minimax M2.1 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Minimax M2.1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Ministral 14B 3.0 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Ministral 14B 3.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Ministral 3B 3.0 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Ministral 3B 3.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Ministral 8B 3.0 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Ministral 8B 3.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Mistral 7B Instruct | Each supported Region: 800 | No | The maximum number of times that you can call InvokeModel in one minute for Mistral mistral-7b-instruct-v0 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Mistral AI Mistral Small | Each supported Region: 400 | No | The maximum number of times that you can call InvokeModel and InvokeModelWithResponseStream in one minute for Mistral AI Mistral Small |
| On-demand model inference requests per minute for Mistral Devstral 2 123b | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Mistral Devstral 2 123b. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Mistral Large | Each supported Region: 400 | No | The maximum number of times that you can call InvokeModel and InvokeModelWithResponseStream in one minute for Mistral mistral-large-2402-v1 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Mistral Large 2407 | Each supported Region: 400 | No | The maximum number of times that you can call model inference in one minute for Mistral Large 2407. The quota considers the combined sum of requests for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream |
| On-demand model inference requests per minute for Mistral Large 3 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Mistral Large 3. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Mistral Mixtral 8x7b Instruct | Each supported Region: 400 | No | The maximum number of times that you can call InvokeModel in one minute for Mistral mixtral-8x7b-v0 |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Moonshot AI Kimi K2.5 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Moonshot AI Kimi K2.5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for NVIDIA Nemotron 3 Super 120B A12B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for NVIDIA Nemotron Nano 2 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for NVIDIA Nemotron Nano 2. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for NVIDIA Nemotron Nano 2 VL | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for NVIDIA Nemotron Nano 2 VL. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Nemotron Nano 3 30B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Nemotron Nano 3 30B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for OpenAI GPT OSS 120B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for OpenAI GPT OSS 120B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for OpenAI GPT OSS 20B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for OpenAI GPT OSS 20B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Qwen3 235B a22b 2507 V1 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 235B a22b 2507 V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Qwen3 32B V1 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 32B V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Qwen3 Coder 30B a3b V1 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 Coder 30B a3b V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Qwen3 Coder 480B a35b V1 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 Coder 480B a35b V1. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Qwen3 Coder Next | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 Coder Next. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Qwen3 Next 80B A3B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 Next 80B A3B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Qwen3 VL 235B A22B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Qwen3 VL 235B A22B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stability.ai Stable Diffusion 3 Large | Each supported Region: 15 | No | The maximum number of times that you can call InvokeModel in one minute for Stability.ai Stable Diffusion 3 Large. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stability.ai Stable Diffusion 3 Medium | Each supported Region: 60 | No | The maximum number of times that you can call InvokeModel in one minute for Stability .ai Stable Diffusion 3 Medium |
| On-demand model inference requests per minute for Stability.ai Stable Diffusion 3.5 Large | Each supported Region: 15 | No | The maximum number of times that you can call InvokeModel in one minute for Stability.ai Stable Diffusion 3.5 Large. |
| On-demand model inference requests per minute for Stability.ai Stable Image Core | Each supported Region: 90 | No | The maximum number of times that you can call InvokeModel in one minute for Stability.ai Stable Image Core. |
| On-demand model inference requests per minute for Stability.ai Stable Image Ultra | Each supported Region: 10 | No | The maximum number of times that you can call InvokeModel in one minute for Stability.ai Stable Image Ultra. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Conservative Upscale | Each supported Region: 2 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Conservative Upscale. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stable Image Control Sketch | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Control Sketch. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Control Structure | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Control Structure. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stable Image Creative Upscale | Each supported Region: 2 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Creative Upscale. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Erase Object | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Erase Object. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stable Image Fast Upscale | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Fast Upscale. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Inpaint | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Inpaint. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stable Image Outpaint | Each supported Region: 2 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Outpaint. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Remove Background | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Remove Background. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stable Image Search and Recolor | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Search and Recolor. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Search and Replace | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Search and Replace. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Stable Image Style Guide | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Style Guide. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Stable Image Style Transfer | Each supported Region: 10 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Stable Image Style Transfer. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Twelve Labs Marengo | Each supported Region: 100 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Twelve Labs Marengo. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Twelve Labs Pegasus | Each supported Region: 60 | Yes | The maximum number of on-demand requests that you can submit for model inference in one minute for Twelve Labs Pegasus. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference requests per minute for TwelveLabs Marengo Embed 3.0 | Each supported Region: 500 | Yes | The maximum number of on-demand requests that you can submit for model inference in one minute for TwelveLabs Marengo Embed 3.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Voxtral Mini 1.0 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Voxtral Mini 1.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Voxtral Small 1.0 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Voxtral Small 1.0. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Writer Palmyra Vision 7B | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Writer Palmyra Vision 7B. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Z.ai GLM 5 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Z.ai GLM 5. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference requests per minute for Z.ai GLM-4.7 | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Z.ai GLM-4.7. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |
| On-demand model inference requests per minute for Z.ai GLM-4.7 Flash | Each supported Region: 10,000 | No | The maximum number of on-demand requests that you can submit for model inference in one minute for Z.ai GLM-4.7 Flash. This quota represents the total number of requests to the Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for AI21 Labs Jamba 1.5 Large | Each supported Region: 300,000 | No | The maximum number of tokens that you can submit for model inference in one minute for AI21 Labs Jamba 1.5 Large. The quota considers the combined sum of input and output tokens across all requests to Converse and InvokeModel. |
| On-demand model inference tokens per minute for AI21 Labs Jamba 1.5 Mini | Each supported Region: 300,000 | No | The maximum number of tokens that you can submit for model inference in one minute for AI21 Labs Jamba 1.5 Mini. The quota considers the combined sum of input and output tokens across all requests to Converse and InvokeModel. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Amazon Nova Lite | us-east-1: 4,000,000<br><br>eu-west-2: 4,000,000<br><br>Each of the other supported Regions: 200,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Amazon Nova Lite. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Amazon Nova Micro | us-east-1: 4,000,000<br><br>eu-west-2: 4,000,000<br><br>Each of the other supported Regions: 200,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Amazon Nova Micro. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Amazon Nova Pro | Each supported Region: 1,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Amazon Nova Pro. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Amazon Titan Image Generator G1 V2 | Each supported Region: 2,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Amazon Titan Image Generator G1 V2. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Amazon Titan Multimodal Embeddings G1 | Each supported Region: 300,000 | No | The maximum number of tokens that you can provide through InvokeModel in one minute for Amazon Titan Multimodal Embeddings G1. |
| On-demand model inference tokens per minute for Amazon Titan Text Embeddings | Each supported Region: 300,000 | No | The maximum number of tokens that you can provide through InvokeModel in one minute for Amazon Titan Text Embeddings. |
| On-demand model inference tokens per minute for Amazon Titan Text Embeddings V2 | Each supported Region: 300,000 | No | The maximum number of tokens that you can provide through InvokeModel in one minute for Amazon Titan Text Embeddings V2. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Amazon Titan Text Express | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Amazon Titan Text Express. |
| On-demand model inference tokens per minute for Amazon Titan Text Premier | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Amazon Titan Text Premier. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Anthropic Claude 3 Haiku | us-east-1: 2,000,000<br><br>us-west-2: 2,000,000<br><br>ap-northeast-1: 200,000<br><br>ap-southeast-1: 200,000<br><br>Each of the other supported Regions: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Haiku. |
| On-demand model inference tokens per minute for Anthropic Claude 3 Opus | Each supported Region: 400,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Anthropic Claude 3 Opus. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Anthropic Claude 3 Sonnet | Each supported Region: 200,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Anthropic Claude 3 Sonnet. |
| On-demand model inference tokens per minute for Anthropic Claude 3.5 Haiku | us-west-1: 300,000<br><br>Each of the other supported Regions: 2,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Anthropic Claude 3.5 Haiku. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Anthropic Claude 3.5 Sonnet | ap-southeast-1: 200,000<br><br>Each of the other supported Regions: 400,000 | No | The maximum number of tokens that you can submit for model inference in one minute for Anthropic Claude 3.5 Sonnet. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Anthropic Claude 3.5 Sonnet V2 | Each supported Region: 400,000 | No | The maximum number of tokens that you can submit for model inference in one minute for Anthropic Claude 3.5 Sonnet V2. The quota considers the combined sum of Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Anthropic Claude 3.7 Sonnet V1 | Each supported Region: 500,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Anthropic Claude 3.7 Sonnet V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Anthropic Claude Opus 4.6 V1 | Each supported Region: 1,500,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Anthropic Claude Opus 4.6 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Anthropic Claude Sonnet 4.6 | Each supported Region: 3,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Anthropic Claude Sonnet 4.6. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Cohere Command R | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Cohere Command R 128k. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Cohere Command R Plus | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Cohere Command R Plus 128k. |
| On-demand model inference tokens per minute for Cohere Embed English | Each supported Region: 300,000 | No | The maximum number of tokens that you can provide through InvokeModel in one minute for Cohere Embed English. |
| On-demand model inference tokens per minute for Cohere Embed Multilingual | Each supported Region: 300,000 | No | The maximum number of tokens that you can provide through InvokeModel in one minute for Cohere Embed Multilingual. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Cohere Embed V4 | Each supported Region: 150,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Cohere Embed V4. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for DeepSeek V3 V1 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for DeepSeek V3 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for DeepSeek V3.2 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for DeepSeek V3.2. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for GPT OSS Safeguard 120B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for GPT OSS Safeguard 120B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for GPT OSS Safeguard 20B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for GPT OSS Safeguard 20B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Gemma 3 12B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Gemma 3 12B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Gemma 3 27B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Gemma 3 27B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Gemma 3 4B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Gemma 3 4B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Kimi K2 Thinking | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Kimi K2 Thinking. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Magistral Small 1.2 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Magistral Small 1.2. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Meta Llama 3 70B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3 70B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Meta Llama 3 8B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3 8B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Meta Llama 3.1 405B Instruct | Each supported Region: 400,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3.1 405B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Meta Llama 3.1 70B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3.1 70B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Meta Llama 3.1 8B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3.1 8B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Meta Llama 3.2 11B Instruct | Each supported Region: 300,000 | No | The maximum number of tokens that you can submit for model inference in one minute for Meta Llama 3.2 11B Instruct. The quota considers the combined sum of tokens for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Meta Llama 3.2 1B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3.2 1B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Meta Llama 3.2 3B Instruct | Each supported Region: 300,000 | No | The maximum number of tokens that you can submit for model inference in one minute for Meta Llama 3.2 3B Instruct. The quota considers the combined sum of tokens for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Meta Llama 3.2 90B Instruct | Each supported Region: 300,000 | No | The maximum number of tokens that you can submit for model inference in one minute for Meta Llama 3.2 90B Instruct. The quota considers the combined sum of tokens for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream. |
| On-demand model inference tokens per minute for Meta Llama 3.3 70B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Meta Llama 3.3 70B Instruct. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for MiniMax M2.5 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for MiniMax M2.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Minimax M2 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Minimax M2. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Minimax M2.1 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Minimax M2.1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Ministral 14B 3.0 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Ministral 14B 3.0. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Ministral 3B 3.0 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Ministral 3B 3.0. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Ministral 8B 3.0 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Ministral 8B 3.0. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Mistral AI Mistral 7B Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Mistral AI Mistral 7B Instruct. |
| On-demand model inference tokens per minute for Mistral AI Mistral Large | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Mistral AI Mistral Large. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Mistral AI Mistral Small | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Mistral AI Mistral Small. |
| On-demand model inference tokens per minute for Mistral AI Mixtral 8X7BB Instruct | Each supported Region: 300,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream for Mistral AI Mixtral 8X7B Instruct. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Mistral Devstral 2 123b | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Mistral Devstral 2 123b. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Mistral Large 2407 | Each supported Region: 300,000 | No | The maximum number of tokens that you can submit for model inference in one minute for Mistral Large 2407. The quota considers the combined sum of tokens for InvokeModel, InvokeModelWithResponseStream, Converse, and ConverseStream |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Mistral Large 3 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Mistral Large 3. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Moonshot AI Kimi K2.5 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Moonshot AI Kimi K2.5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for NVIDIA Nemotron 3 Super 120B A12B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for NVIDIA Nemotron Nano 2 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for NVIDIA Nemotron Nano 2. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for NVIDIA Nemotron Nano 2 VL | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for NVIDIA Nemotron Nano 2 VL. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Nemotron Nano 3 30B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Nemotron Nano 3 30B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for OpenAI GPT OSS 120B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for OpenAI GPT OSS 120B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for OpenAI GPT OSS 20B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for OpenAI GPT OSS 20B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Qwen3 235B a22b 2507 V1 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 235B a22b 2507 V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Qwen3 32B V1 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 32B V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Qwen3 Coder 30B a3b V1 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 Coder 30B a3b V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Qwen3 Coder 480B a35b V1 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 Coder 480B a35b V1. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Qwen3 Coder Next | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 Coder Next. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Qwen3 Next 80B A3B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 Next 80B A3B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Qwen3 VL 235B A22B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Qwen3 VL 235B A22B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Voxtral Mini 1.0 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Voxtral Mini 1.0. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Voxtral Small 1.0 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Voxtral Small 1.0. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Writer Palmyra Vision 7B | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Writer Palmyra Vision 7B. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Z.ai GLM 5 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Z.ai GLM 5. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| On-demand model inference tokens per minute for Z.ai GLM-4.7 | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Z.ai GLM-4.7. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| On-demand model inference tokens per minute for Z.ai GLM-4.7 Flash | Each supported Region: 100,000,000 | No | The maximum number of on-demand tokens that you can submit for model inference in one minute for Z.ai GLM-4.7 Flash. The quota considers the combined sum of input and output tokens across all requests to Converse, ConverseStream, InvokeModel and InvokeModelWithResponseStream. |
| Parameters per function | Each supported Region: 5 | Yes | The maximum number of parameters that you can have in an action group function. |
| PrepareAgent requests per second | Each supported Region: 2 | No | The maximum number of PrepareAgent API requests per second. |
| Records per batch inference job for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Amazon Nova 2 Multimodal Embeddings V1. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Amazon Nova Premier | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Amazon Nova Premier. |
| Records per batch inference job for Claude 3 Haiku | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3 Haiku. |
| Records per batch inference job for Claude 3 Opus | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3 Opus. |
| Records per batch inference job for Claude 3 Sonnet | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3 Sonnet. |
| Records per batch inference job for Claude 3.5 Haiku | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3.5 Haiku. |
| Records per batch inference job for Claude 3.5 Sonnet | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3.5 Sonnet. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Claude 3.5 Sonnet v2 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3.5 Sonnet v2. |
| Records per batch inference job for Claude 3.7 Sonnet | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude 3.7 Sonnet. |
| Records per batch inference job for Claude Haiku 4.5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Haiku 4.5. |
| Records per batch inference job for Claude Opus 4.5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Opus 4.5. |
| Records per batch inference job for Claude Opus 4.6 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Opus 4.6. |
| Records per batch inference job for Claude Opus 5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Opus 5. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Records per batch inference job for Claude Sonnet 4 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Sonnet 4. |
| Records per batch inference job for Claude Sonnet 4.5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Sonnet 4.5. |
| Records per batch inference job for Claude Sonnet 4.6 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Claude Sonnet 4.6. |
| Records per batch inference job for DeepSeek V3.2 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for DeepSeek V3.2. |
| Records per batch inference job for DeepSeek v3 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for DeepSeek v3. |
| Records per batch inference job for Devstral 2 123B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Devstral 2 123B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for GLM 4.7 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for GLM 4.7. |
| Records per batch inference job for GLM 4.7 Flash | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for GLM 4.7 Flash. |
| Records per batch inference job for GLM 5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for GLM 5. |
| Records per batch inference job for Gemma 3 12B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Gemma 3 12B. |
| Records per batch inference job for Gemma 3 27B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Gemma 3 27B. |
| Records per batch inference job for Gemma 3 4B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Gemma 3 4B. |
| Records per batch inference job for Kimi K2 Thinking | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Kimi K2 Thinking. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Kimi K2.5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Kimi K2.5. |
| Records per batch inference job for Llama 3.1 405B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.1 405B Instruct. |
| Records per batch inference job for Llama 3.1 70B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.1 70B Instruct. |
| Records per batch inference job for Llama 3.1 8B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.1 8B Instruct. |
| Records per batch inference job for Llama 3.2 11B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.2 11B Instruct. |
| Records per batch inference job for Llama 3.2 1B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job Llama 3.2 1B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Llama 3.2 3B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.2 3B Instruct. |
| Records per batch inference job for Llama 3.2 90B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.2 90B Instruct. |
| Records per batch inference job for Llama 3.3 70B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 3.3 70B Instruct. |
| Records per batch inference job for Llama 4 Maverick | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 4 Maverick. |
| Records per batch inference job for Llama 4 Scout | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Llama 4 Scout. |
| Records per batch inference job for Magistral Small 2509 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Magistral Small 2509. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for MiniMax M2 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for MiniMax M2. |
| Records per batch inference job for MiniMax M2.1 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for MiniMax M2.1. |
| Records per batch inference job for MiniMax M2.5 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for MiniMax M2.5. |
| Records per batch inference job for Ministral 3 14B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Ministral 3 14B. |
| Records per batch inference job for Ministral 3 8B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Ministral 3 8B. |
| Records per batch inference job for Ministral 3B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Ministral 3B. |
| Records per batch inference job for Mistral Large 2 (24.07) | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Mistral Large 2 (24.07). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Mistral Large 3 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Mistral Large 3. |
| Records per batch inference job for Mistral Small | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Mistral Small. |
| Records per batch inference job for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for NVIDIA Nemotron 3 Super 120B A12B. |
| Records per batch inference job for NVIDIA Nemotron Nano 12B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for NVIDIA Nemotron Nano 12B. |
| Records per batch inference job for NVIDIA Nemotron Nano 3 30B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for NVIDIA Nemotron Nano 3 30B. |
| Records per batch inference job for NVIDIA Nemotron Nano 9B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for NVIDIA Nemotron Nano 9B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Nova 2 Lite | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Nova 2 Lite. |
| Records per batch inference job for Nova Lite V1 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Nova Lite V1. |
| Records per batch inference job for Nova Micro V1 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Nova Micro V1. |
| Records per batch inference job for Nova Pro V1 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Nova Pro V1. |
| Records per batch inference job for OpenAI GPT OSS 120b | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for OpenAI GPT OSS 120b. |
| Records per batch inference job for OpenAI GPT OSS 20b | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for OpenAI GPT OSS 20b. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for OpenAI GPT OSS Safeguard 120b | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for OpenAI GPT OSS Safeguard 120b. |
| Records per batch inference job for OpenAI GPT OSS Safeguard 20b | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for OpenAI GPT OSS Safeguard 20b. |
| Records per batch inference job for Qwen3 235B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 235B. |
| Records per batch inference job for Qwen3 32B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 32B. |
| Records per batch inference job for Qwen3 Coder 30B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 Coder 30B. |
| Records per batch inference job for Qwen3 Coder 480B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 Coder 480B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per batch inference job for Qwen3 Coder Next | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 Coder Next. |
| Records per batch inference job for Qwen3 Next 80B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 Next 80B. |
| Records per batch inference job for Qwen3 VL 235B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Qwen3 VL 235B. |
| Records per batch inference job for Titan Multimodal Embeddings G1 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Titan Multimodal Embeddings G1. |
| Records per batch inference job for Titan Text Embeddings V2 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Titan Text Embeddings V2. |
| Records per batch inference job for Voxtral Mini 3B 2507 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Voxtral Mini 3B 2507. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Records per batch inference job for Voxtral Small 24B 2507 | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Voxtral Small 24B 2507. |
| Records per batch inference job for Writer Palmyra Vision 7B | Each supported Region: 100,000 | Yes | The maximum number of records across all input files in a batch inference job for Writer Palmyra Vision 7B. |
| Records per input file per batch inference job for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Amazon Nova 2 Multimodal Embeddings V1. |
| Records per input file per batch inference job for Amazon Nova Premier | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Amazon Nova Premier. |
| Records per input file per batch inference job for Claude 3 Haiku | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3 Haiku. |
| Records per input file per batch inference job for Claude 3 Opus | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3 Opus. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Claude 3 Sonnet | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3 Sonnet. |
| Records per input file per batch inference job for Claude 3.5 Haiku | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3.5 Haiku. |
| Records per input file per batch inference job for Claude 3.5 Sonnet | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3.5 Sonnet. |
| Records per input file per batch inference job for Claude 3.5 Sonnet v2 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3.5 Sonnet v2. |
| Records per input file per batch inference job for Claude 3.7 Sonnet | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude 3.7 Sonnet. |
| Records per input file per batch inference job for Claude Haiku 4.5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Haiku 4.5. |
| Records per input file per batch inference job for Claude Opus 4.5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Opus 4.5. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Claude Opus 4.6 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Opus 4.6. |
| Records per input file per batch inference job for Claude Opus 5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Opus 5. |
| Records per input file per batch inference job for Claude Sonnet 4 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Sonnet 4. |
| Records per input file per batch inference job for Claude Sonnet 4.5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Sonnet 4.5. |
| Records per input file per batch inference job for Claude Sonnet 4.6 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Claude Sonnet 4.6. |
| Records per input file per batch inference job for DeepSeek V3.2 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for DeepSeek V3.2. |
| Records per input file per batch inference job for DeepSeek v3 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for DeepSeek v3. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Devstral 2 123B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Devstral 2 123B. |
| Records per input file per batch inference job for GLM 4.7 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for GLM 4.7. |
| Records per input file per batch inference job for GLM 4.7 Flash | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for GLM 4.7 Flash. |
| Records per input file per batch inference job for GLM 5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for GLM 5. |
| Records per input file per batch inference job for Gemma 3 12B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Gemma 3 12B. |
| Records per input file per batch inference job for Gemma 3 27B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Gemma 3 27B. |
| Records per input file per batch inference job for Gemma 3 4B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Gemma 3 4B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Kimi K2 Thinking | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Kimi K2 Thinking. |
| Records per input file per batch inference job for Kimi K2.5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Kimi K2.5. |
| Records per input file per batch inference job for Llama 3.1 405B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.1 405B Instruct. |
| Records per input file per batch inference job for Llama 3.1 70B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.1 70B Instruct. |
| Records per input file per batch inference job for Llama 3.1 8B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.1 8B Instruct. |
| Records per input file per batch inference job for Llama 3.2 11B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.2 11B Instruct. |
| Records per input file per batch inference job for Llama 3.2 1B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job Llama 3.2 1B Instruct. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Llama 3.2 3B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.2 3B Instruct. |
| Records per input file per batch inference job for Llama 3.2 90B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.2 90B Instruct. |
| Records per input file per batch inference job for Llama 3.3 70B Instruct | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 3.3 70B Instruct. |
| Records per input file per batch inference job for Llama 4 Maverick | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 4 Maverick. |
| Records per input file per batch inference job for Llama 4 Scout | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Llama 4 Scout. |
| Records per input file per batch inference job for Magistral Small 2509 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Magistral Small 2509. |
| Records per input file per batch inference job for MiniMax M2 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for MiniMax M2. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for MiniMax M2.1 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for MiniMax M2.1. |
| Records per input file per batch inference job for MiniMax M2.5 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for MiniMax M2.5. |
| Records per input file per batch inference job for Ministral 3 14B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Ministral 3 14B. |
| Records per input file per batch inference job for Ministral 3 8B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Ministral 3 8B. |
| Records per input file per batch inference job for Ministral 3B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Ministral 3B. |
| Records per input file per batch inference job for Mistral Large 2 (24.07) | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Mistral Large 2 (24.07). |
| Records per input file per batch inference job for Mistral Large 3 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Mistral Large 3. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Mistral Small | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Mistral Small. |
| Records per input file per batch inference job for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for NVIDIA Nemotron 3 Super 120B A12B. |
| Records per input file per batch inference job for NVIDIA Nemotron Nano 12B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for NVIDIA Nemotron Nano 12B. |
| Records per input file per batch inference job for NVIDIA Nemotron Nano 3 30B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for NVIDIA Nemotron Nano 3 30B. |
| Records per input file per batch inference job for NVIDIA Nemotron Nano 9B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for NVIDIA Nemotron Nano 9B. |
| Records per input file per batch inference job for Nova 2 Lite | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Nova 2 Lite. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Nova Lite V1 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Nova Lite V1. |
| Records per input file per batch inference job for Nova Micro V1 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Nova Micro V1. |
| Records per input file per batch inference job for Nova Pro V1 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Nova Pro V1. |
| Records per input file per batch inference job for OpenAI GPT OSS 120b | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for OpenAI GPT OSS 120b. |
| Records per input file per batch inference job for OpenAI GPT OSS 20b | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for OpenAI GPT OSS 20b. |
| Records per input file per batch inference job for OpenAI GPT OSS Safeguard 120b | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for OpenAI GPT OSS Safeguard 120b. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for OpenAI GPT OSS Safeguard 20b | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for OpenAI GPT OSS Safeguard 20b. |
| Records per input file per batch inference job for Qwen3 235B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 235B. |
| Records per input file per batch inference job for Qwen3 32B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 32B. |
| Records per input file per batch inference job for Qwen3 Coder 30B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 Coder 30B. |
| Records per input file per batch inference job for Qwen3 Coder 480B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 Coder 480B. |
| Records per input file per batch inference job for Qwen3 Coder Next | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 Coder Next. |
| Records per input file per batch inference job for Qwen3 Next 80B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 Next 80B. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Records per input file per batch inference job for Qwen3 VL 235B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Qwen3 VL 235B. |
| Records per input file per batch inference job for Titan Multimodal Embeddings G1 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Titan Multimodal Embeddings G1. |
| Records per input file per batch inference job for Titan Text Embeddings V2 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Titan Text Embeddings V2. |
| Records per input file per batch inference job for Voxtral Mini 3B 2507 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Voxtral Mini 3B 2507. |
| Records per input file per batch inference job for Voxtral Small 24B 2507 | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Voxtral Small 24B 2507. |
| Records per input file per batch inference job for Writer Palmyra Vision 7B | Each supported Region: 100,000 | Yes | The maximum number of records in an input file in a batch inference job for Writer Palmyra Vision 7B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Amazon Nova 2 Multimodal Embeddings V1 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Amazon Nova 2 Multimodal Embeddings V1. |
| Sum of in-progress and submitted batch inference jobs using a base model for Amazon Nova Premier | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Amazon Nova Premier. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3 Haiku | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3 Haiku. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3 Opus | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3 Opus. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3 Sonnet | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3 Sonnet. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3.5 Haiku | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3.5 Haiku. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3.5 Sonnet | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3.5 Sonnet. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3.5 Sonnet v2 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3.5 Sonnet v2. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude 3.7 Sonnet | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude 3.7 Sonnet. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Haiku 4.5 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Haiku 4.5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Opus 4.5 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Opus 4.5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Opus 4.6 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Opus 4.6. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Opus 5 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Opus 5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Sonnet 4 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Sonnet 4. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Sonnet 4.5. | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Sonnet 4.5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Claude Sonnet 4.6 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Claude Sonnet 4.6. |
| Sum of in-progress and submitted batch inference jobs using a base model for DeepSeek V3.2 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for DeepSeek V3.2. |
| Sum of in-progress and submitted batch inference jobs using a base model for DeepSeek v3 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for DeepSeek v3. |
| Sum of in-progress and submitted batch inference jobs using a base model for Devstral 2 123B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Devstral 2 123B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for GLM 4.7 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for GLM 4.7. |
| Sum of in-progress and submitted batch inference jobs using a base model for GLM 4.7 Flash | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for GLM 4.7 Flash. |
| Sum of in-progress and submitted batch inference jobs using a base model for GLM 5 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for GLM 5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Gemma 3 12B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Gemma 3 12B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Gemma 3 27B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Gemma 3 27B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Gemma 3 4B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Gemma 3 4B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Kimi K2 Thinking | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Kimi K2 Thinking. |
| Sum of in-progress and submitted batch inference jobs using a base model for Kimi K2.5 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Kimi K2.5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.1 405B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.1 405B Instruct. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.1 70B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.1 70B Instruct. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.1 8B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.1 8B Instruct. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.2 11B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.2 11B Instruct. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.2 1B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.2 1B Instruct. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.2 3B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.2 3B Instruct. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.2 90B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.2 90B Instruct. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 3.3 70B Instruct | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 3.3 70B Instruct. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 4 Maverick | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 4 Maverick. |
| Sum of in-progress and submitted batch inference jobs using a base model for Llama 4 Scout | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Llama 4 Scout. |
| Sum of in-progress and submitted batch inference jobs using a base model for Magistral Small 2509 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Magistral Small 2509. |
| Sum of in-progress and submitted batch inference jobs using a base model for MiniMax M2 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for MiniMax M2. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for MiniMax M2.1 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for MiniMax M2.1. |
| Sum of in-progress and submitted batch inference jobs using a base model for MiniMax M2.5 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for MiniMax M2.5. |
| Sum of in-progress and submitted batch inference jobs using a base model for Ministral 3 14B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Ministral 3 14B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Ministral 3 8B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Ministral 3 8B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Ministral 3B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Ministral 3B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Mistral Large 2 (24.07) | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Mistral Large 2 (24.07). |
| Sum of in-progress and submitted batch inference jobs using a base model for Mistral Large 3 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Mistral Large 3. |
| Sum of in-progress and submitted batch inference jobs using a base model for Mistral Small | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Mistral Small. |
| Sum of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron 3 Super 120B A12B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron 3 Super 120B A12B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron Nano 12B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron Nano 12B. |
| Sum of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron Nano 3 30B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron Nano 3 30B. |
| Sum of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron Nano 9B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for NVIDIA Nemotron Nano 9B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Nova 2 Lite | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Nova 2 Lite. |
| Sum of in-progress and submitted batch inference jobs using a base model for Nova Lite V1 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Nova Lite V1. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Nova Micro V1 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Nova Micro V1. |
| Sum of in-progress and submitted batch inference jobs using a base model for Nova Pro V1 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Nova Pro V1. |
| Sum of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS 120b | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS 120b. |
| Sum of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS 20b | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS 20b. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Sum of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS Safeguard 120b | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS Safeguard 120b. |
| Sum of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS Safeguard 20b | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for OpenAI GPT OSS Safeguard 20b. |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 235B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 235B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 32B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 32B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 Coder 30B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 Coder 30B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 Coder 480B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 Coder 480B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 Coder Next | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 Coder Next. |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 Next 80B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 Next 80B. |
| Sum of in-progress and submitted batch inference jobs using a base model for Qwen3 VL 235B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Qwen3 VL 235B. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Titan Multimodal Embeddings G1 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Titan Multimodal Embeddings G1. |
| Sum of in-progress and submitted batch inference jobs using a base model for Titan Text Embeddings V2 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Titan Text Embeddings V2. |
| Sum of in-progress and submitted batch inference jobs using a base model for Voxtral Mini 3B 2507 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Voxtral Mini 3B 2507. |
| Sum of in-progress and submitted batch inference jobs using a base model for Voxtral Small 24B 2507 | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Voxtral Small 24B 2507. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sum of in-progress and submitted batch inference jobs using a base model for Writer Palmyra Vision 7B | Each supported Region: 100 | Yes | The maximum number of in-progress and submitted batch inference jobs using a base model for Writer Palmyra Vision 7B. |
| Sum of in-progress and submitted batch inference jobs using a custom model for Titan Multimodal Embeddings G1 | Each supported Region: 3 | No | The maximum number of in-progress and submitted batch inference jobs using a custom model for Titan Multimodal Embeddings G1. |
| Sum of in-progress and submitted batch inference jobs using a custom model for Titan Text Embeddings V2 | Each supported Region: 3 | No | The maximum number of in-progress and submitted batch inference jobs using a custom model for Titan Text Embeddings V2 |
| Throttle rate limit for Bedrock Data Automation Runtime: ListTagsForResource | Each supported Region: 25 per second | No | The maximum number of Bedrock Data Automation Runtime: ListTagsForResource requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for Bedrock Data Automation Runtime: TagResource | Each supported Region: 25 per second | No | The maximum number of Bedrock Data Automation Runtime: TagResource requests you can make per second per account, in the current region |
| Throttle rate limit for Bedrock Data Automation Runtime: UntagResource | Each supported Region: 25 per second | No | The maximum number of Bedrock Data Automation Runtime: UntagResource requests you can make per second per account, in the current region |
| Throttle rate limit for Bedrock Data Automation: ListTagsForResource | Each supported Region: 25 per second | No | The maximum number of Bedrock Data Automation: ListTagsForResource requests you can make per second per account, in the current region |
| Throttle rate limit for Bedrock Data Automation: TagResource | Each supported Region: 25 per second | No | The maximum number of Bedrock Data Automation: TagResource requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for Bedrock Data Automation: UntagResource | Each supported Region: 25 per second | No | The maximum number of Bedrock Data Automation: UntagResource requests you can make per second per account, in the current region |
| Throttle rate limit for CreateBlueprint | Each supported Region: 5 per second | No | The maximum number of CreateBlueprint requests you can make per second per account, in the current region |
| Throttle rate limit for CreateBlueprintVersion | Each supported Region: 5 per second | No | The maximum number of CreateBlueprintVersion requests you can make per second per account, in the current region |
| Throttle rate limit for CreateDataAutomationLibrary | Each supported Region: 3 per second | No | The maximum number of CreateDataAutomationLibrary requests you can make per second per account, in the current region |
| Throttle rate limit for CreateDataAutomationProject | Each supported Region: 5 per second | No | The maximum number of CreateDataAutomationProject requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate limit for DeleteBlueprint | Each supported Region: 5 per second | No | The maximum number of DeleteBlueprint requests you can make per second per account, in the current region |
| Throttle rate limit for DeleteDataAutomationLibrary | Each supported Region: 3 per second | No | The maximum number of DeleteDataAutomationLibrary requests you can make per second per account, in the current region |
| Throttle rate limit for DeleteDataAutomationProject | Each supported Region: 5 per second | No | The maximum number of DeleteDataAutomationProject requests you can make per second per account, in the current region |
| Throttle rate limit for GetBlueprint | Each supported Region: 5 per second | No | The maximum number of GetBlueprint requests you can make per second per account, in the current region |
| Throttle rate limit for GetDataAutomationLibrary | Each supported Region: 5 per second | No | The maximum number of GetDataAutomationLibrary requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate limit for GetDataAutomationLibraryEntity | Each supported Region: 5 per second | No | The maximum number of GetDataAutomationLibraryEntity requests you can make per second per account, in the current region |
| Throttle rate limit for GetDataAutomationLibraryIngestionJob | Each supported Region: 5 per second | No | The maximum number of GetDataAutomationLibraryIngestionJob requests you can make per second per account, in the current region |
| Throttle rate limit for GetDataAutomationProject | Each supported Region: 5 per second | No | The maximum number of GetDataAutomationProject requests you can make per second per account, in the current region |
| Throttle rate limit for GetDataAutomationStatus | Each supported Region: 10 per second | No | The maximum number of GetDataAutomationStatus requests you can make per second per account, in the current region |
| Throttle rate limit for InvokeDataAutomationAsync | Each supported Region: 10 per second | No | The maximum number of InvokeDataAutomationAsync requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate limit for InvokeDataAutomationLibraryIngestionJob | Each supported Region: 5 per second | No | The maximum number of InvokeDataAutomationLibraryIngestionJob requests you can make per second per account, in the current region |
| Throttle rate limit for ListBlueprints | Each supported Region: 5 per second | No | The maximum number of ListBlueprints requests you can make per second per account, in the current region |
| Throttle rate limit for ListDataAutomationLibraries | Each supported Region: 5 per second | No | The maximum number of ListDataAutomationLibraries requests you can make per second per account, in the current region |
| Throttle rate limit for ListDataAutomationLibraryEntities | Each supported Region: 5 per second | No | The maximum number of ListDataAutomationLibraryEntities requests you can make per second per account, in the current region |
| Throttle rate limit for ListDataAutomationLibraryIngestionJobs | Each supported Region: 5 per second | No | The maximum number of ListDataAutomationLibraryIngestionJobs requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for ListDataAutomationProjects | Each supported Region: 5 per second | No | The maximum number of ListDataAutomationProjects requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateBlueprint | Each supported Region: 5 per second | No | The maximum number of UpdateBlueprint requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateDataAutomationLibrary | Each supported Region: 5 per second | No | The maximum number of UpdateDataAutomationLibrary requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateDataAutomationProject | Each supported Region: 5 per second | No | The maximum number of UpdateDataAutomationProject requests you can make per second per account, in the current region |
| UpdateAgent requests per second | Each supported Region: 4 | No | The maximum number of UpdateAgent API requests per second. |
| UpdateAgentActionGroup requests per second | Each supported Region: 6 | No | The maximum number of UpdateAgentActionGroup API requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateAgentAlias requests per second | Each supported Region: 2 | No | The maximum number of UpdateAgentAlias API requests per second. |
| UpdateAgentKnowledgeBase requests per second | Each supported Region: 4 | No | The maximum number of UpdateAgentKnowledgeBase API requests per second. |
| [bedrock-mantle endpoint] Input tokens per minute for Claude Fable 5 | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to Claude Fable 5 in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for Claude Fable 5.1 | Each supported Region: 5,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to Claude Fable 5.1 in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Input tokens per minute for Claude Opus 4.7 | Each supported Region: 20,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to Claude Opus 4.7 in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for Claude Opus 4.8 | Each supported Region: 20,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to Claude Opus 4.8 in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for Claude Opus 5 | Each supported Region: 20,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to Claude Opus 5 in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Input tokens per minute for Claude Sonnet 5 | Each supported Region: 3,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to Claude Sonnet 5 in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for GPT-5.4 | Each supported Region: 20,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to GPT-5.4 in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for GPT-5.5 | Each supported Region: 10,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to GPT-5.5 in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for GPT-5.6 Luna | Each supported Region: 20,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to GPT-5.6 Luna in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Input tokens per minute for GPT-5.6 Sol | Each supported Region: 10,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to GPT-5.6 Sol in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Input tokens per minute for GPT-5.6 Terra | Each supported Region: 20,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of input tokens per minute your account can submit to GPT-5.6 Terra in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for Claude Fable 5 | Each supported Region: 200,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute Claude Fable 5 can generate for your account in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Output tokens per minute for Claude Fable 5.1 | Each supported Region: 500,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute Claude Fable 5.1 can generate for your account in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for Claude Opus 4.7 | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute Claude Opus 4.7 can generate for your account in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for Claude Opus 4.8 | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute Claude Opus 4.8 can generate for your account in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Output tokens per minute for Claude Opus 5 | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute Claude Opus 5 can generate for your account in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for Claude Sonnet 5 | Each supported Region: 300,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute Claude Sonnet 5 can generate for your account in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for GPT-5.4 | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute GPT-5.4 can generate for your account in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Output tokens per minute for GPT-5.5 | Each supported Region: 1,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute GPT-5.5 can generate for your account in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for GPT-5.6 Luna | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute GPT-5.6 Luna can generate for your account in this region. Shared across all APIs for this model. |
| [bedrock-mantle endpoint] Output tokens per minute for GPT-5.6 Sol | Each supported Region: 1,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute GPT-5.6 Sol can generate for your account in this region. Shared across all APIs for this model. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| [bedrock-mantle endpoint] Output tokens per minute for GPT-5.6 Terra | Each supported Region: 2,000,000 | Yes | [bedrock-mantle endpoint] The maximum number of output tokens per minute GPT-5.6 Terra can generate for your account in this region. Shared across all APIs for this model. |

**Note**

You can request a quota increase for your account by following the steps below:

- If a quota is marked as **Yes**, you can adjust it by following the steps at Requesting a Quota Increase in the Service Quotas User Guide.
- If a quota is marked as **No**, you can submit a request through the limit increase form to be considered for an increase.
- For any model, you can request an increase for the following quotas together:
  - Cross-Region InvokeModel tokens per minute for *${model}*
  - Cross-Region InvokeModel requests per minute for *${model}*
  - On-demand InvokeModel tokens per minute for *${model}*
  - On-demand InvokeModel requests per minute for *${model}*

  To request an increase for any combination of these quotas, request an increase for the **Cross-Region InvokeModel tokens per minute for *${model}*** quota by following the steps at Requesting a Quota Increase in the Service Quotas User Guide. After you do so, the support team will reach out and offer you the option of also increasing the other three quotas. Due to overwhelming demand, priority will be given to customers who generate traffic that consumes their existing quota allocation. Your request might be denied if you don't meet this condition.

# Amazon Bedrock AgentCore endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Amazon Bedrock AgentCore control plane APIs

The following table provides a list of Region-specific endpoints that Amazon Bedrock AgentCore supports for calling control plane operations.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | bedrock-agentcore-control.us-east-1.amazonaws.com | HTTPS |
| US East (Ohio) | us-east-2 | bedrock-agentcore-control.us-east-2.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | bedrock-agentcore-control.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | bedrock-agentcore-control. ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock-agentcore-control.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock-agentcore-control.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock-agentcore-control.ap-northeast-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | bedrock-agentcore-control.eu-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | bedrock-agentcore-control.eu-central-1.amazonaws.com | HTTPS |

## Amazon Bedrock AgentCore data plane APIs

The following table provides a list of Region-specific endpoints that Amazon Bedrock AgentCore supports for calling data plane operations.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | bedrock-agentcore.us-east-1.amazonaws.com | HTTPS |
| US East (Ohio) | us-east-2 | bedrock-agentcore.us-east-2.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | bedrock-agentcore.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | bedrock-agentcore.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | bedrock-agentcore.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | bedrock-agentcore.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | bedrock-agentcore.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | bedrock-agentcore.eu-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | bedrock-agentcore.eu-central-1.amazonaws.com | HTTPS |

## Amazon Bedrock AgentCore Gateway data plane APIs

The following table provides a list of Region-specific endpoints that Amazon Bedrock AgentCore Gateway supports for invoking a gateway.

Replace *gatewayId* with the ID of your gateway.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | *gatewayId* .gateway.bedrock-agentcore.us-east-1.amazonaws.com | HTTPS |
| US East (Ohio) | us-east-2 | *gatewayId* .gateway.bedrock-agentcore.us-east-2.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | *gatewayId* .gateway.bedrock-agentcore.us-west-2.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | *gatewayId* .gateway.bedrock-agentcore.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | *gatewayId* .gateway.bedrock-agentcore.eu-west-1.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | *gatewayId* .gateway.bedrock-agentcore.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | *gatewayId* .gateway.bedrock-agentcore.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | *gatewayId* .gateway.bedrock-agentcore.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | *gatewayId* .gateway.bedrock-agentcore.ap-southeast-2.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | *gatewayId* .gateway.bedrock-agentcore.us-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | *gatewayId* .gateway.bedrock-agentcore.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | *gatewayId* .gateway.bedrock-agentcore.eu-west-3.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | *gatewayId* .gateway.bedrock-agentcore.eu-south-1.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | *gatewayId* .gateway.bedrock-agentcore.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | *gatewayId* .gateway.bedrock-agentcore.eu-north-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | *gatewayId* .gateway.bedrock-agentcore.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | *gatewayId* .gateway.bedrock-agentcore.ap-south-2.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | *gatewayId* .gateway.bedrock-agentcore.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | *gatewayId* .gateway.bedrock-agentcore.ap-southeast-7.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | *gatewayId* .gateway.bedrock-agentcore.ca-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | *gatewayId* .gateway.bedrock-agentcore.sa-east-1.amazonaws.com | HTTPS |

### Service quotas

For information about Amazon Bedrock AgentCore service quotas, see Quotas for Amazon Bedrock AgentCore.

## AWS Billing and Cost Management endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

AWS Billing and Cost Management includes the AWS Cost Explorer API, the AWS Cost and Usage Reports API, the AWS Budgets API, and the AWS Price List API.

# Service endpoints

## AWS Cost Explorer

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ce.us-east-1.amazonaws.com | HTTPS |
| | | ce.us-east-1.api.aws | HTTPS |

## AWS Budgets

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | budgets.amazonaws.com | HTTPS |
| | | budgets.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | budgets.amazonaws.com | HTTPS |
| | | budgets.us-east-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | budgets.amazonaws.com | HTTPS |
| | | budgets.us-east-1.api.aws | HTTPS |

## AWS Cost Optimization Hub

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | cost-optimization-hub.us-east-1.amazonaws.com<br><br>cost-optimization-hub.us-east-1.api.aws | HTTPS<br><br>HTTPS | |

## AWS Data Exports

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | bcm-data-exports.us-east-1.api.aws | | |

## AWS Free Tier

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | freetier.us-east-1.api.aws | | |

## AWS Price List Service

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | api.pricing.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | api.pricing.ap-south-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | api.pricing.eu-central-1.amazonaws.com | HTTPS | |

## Savings Plans

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | savingsplans.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | savingsplans.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | savingsplans.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | savingsplans.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | savingsplans.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | savingsplans.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | savingsplans.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | savingsplans.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | savingsplans.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | savingsplans.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | savingsplans.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | savingsplans.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | savingsplans.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | savingsplans.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | savingsplans.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | savingsplans.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | savingsplans.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Cost Explorer saved reports | Each supported Region: 300 | No | The maximum number of reports that you save in an AWS account. |
| Number of AWS managed linked account, cost category, and cost allocation tag anomaly monitors | Each supported Region: 1 | No | Maximum number of AWS managed anomaly monitors (linked account, cost category, or cost allocation tags) per account |
| Number of AWS managed services anomaly monitors | Each supported Region: 1 | No | The maximum number of AWS managed service anomaly monitors per account |
| Number of Cost Categories | Each supported Region: 50 | No | The maximum number of Cost Categories per payer account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of active cost allocation tag keys | Each supported Region: 500 | Yes | The maximum number of active cost allocation tag keys per payer account |
| Number of affected accounts in an organization that can be migrated using Bulk Policy Migrator | Each supported Region: 200 | Yes | The maximum number of affected accounts in an organization that you can migrate with Bulk Policy Migrator |
| Number of affected policies in an organization that can be migrated using Bulk Policy Migrator | Each supported Region: 1,000 | Yes | The maximum number of affected policies in an organization that you can migrate with Bulk Policy Migrator |
| Number of anomaly alert subscriptions | Each supported Region: 100 | No | The maximum number of anomaly alert subscriptions you can create per account |
| Number of customer managed anomaly monitors | Each supported Region: 500 | No | Maximum number of customer managed anomaly monitors (linked account, cost category, or cost allocation tags) per account |
| Number of refresh requests for Savings Plans recommendations per day | Each supported Region: 3 per 24 hours | No | The maximum number of daily refresh requests for Savings Plans recommendations per consolidated billing family |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Number of rules per Cost Category (API) | Each supported Region: 500 | No | The maximum number of Cost Category rules per Cost Category (API) |
| Number of rules per Cost Category (UI) | Each supported Region: 100 | No | The maximum number of Cost Category rules per Cost Category (UI) |
| Number of values tracked by AWS managed anomaly monitors | Each supported Region: 5,000 | No | Maximum tracked dimension values per AWS managed anomaly monitor |

For more information, see AWS Billing quotas and restrictions and AWS Cost Management quotas and restrictions.

## AWS Billing Conductor endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | billingconductor.us-east-1.amazonaws.com | HTTPS | |
| | | billingconductor.us-east-1.api.aws | HTTPS | |

## Service quotas

For more information about AWS Billing Conductor service quotas and restrictions, see Quotas and restrictions in the *AWS Billing Conductor User Guide*.

# Amazon Braket endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | braket.us-east-1.amazonaws.com | HTTPS | |
| | | braket-fips.us-east-1.api.aws | HTTPS | |
| | | braket.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | braket.us-west-1.amazonaws.com | HTTPS |
| | | braket-fips.us-west-1.api.aws | HTTPS |
| | | braket.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | braket.us-west-2.amazonaws.com | HTTPS |
| | | braket-fips.us-west-2.api.aws | HTTPS |
| | | braket.us-west-2.api.aws | HTTPS |
| Europe (London) | eu-west-2 | braket.eu-west-2.amazonaws.com | HTTPS |
| | | braket-fips.eu-west-2.api.aws | HTTPS |
| | | braket.eu-west-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | braket.eu-north-1.amazonaws.com | HTTPS |
| | | braket-fips.eu-north-1.api.aws | HTTPS |
| | | braket.eu-north-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Burst rate of API requests | Each supported Region: 600 | No | The maximum number of additional requests per second (RPS) that you can send in one burst in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst rate of CancelJob requests | Each supported Region: 5 | No | The maximum number of additional CancelJob requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of CancelQuantumTask requests | Each supported Region: 20 | No | The maximum number of additional CancelQuantumTask requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of CreateJob requests | Each supported Region: 5 | No | The maximum number of additional CreateJob requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of CreateQuantumTask requests | Each supported Region: 40 | No | The maximum number of additional CreateQuantumTask requests per second (RPS) that you can send in one burst in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst rate of CreateSpendingLimit requests | Each supported Region: 5 | No | The maximum number of additional CreateSpendingLimit requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of DeleteSpendingLimit requests | Each supported Region: 5 | No | The maximum number of additional DeleteSpendingLimit requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of GetDevice requests | Each supported Region: 50 | No | The maximum number of additional GetDevice requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of GetJob requests | Each supported Region: 25 | No | The maximum number of additional GetJob requests per second (RPS) that you can send in one burst in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst rate of GetQuantumTask requests | Each supported Region: 500 | No | The maximum number of additional GetQuantumTask requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of SearchDevices requests | Each supported Region: 50 | No | The maximum number of additional SearchDevices requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of SearchJobs requests | Each supported Region: 50 | No | The maximum number of additional SearchJobs requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of SearchQuantumTasks requests | Each supported Region: 50 | No | The maximum number of additional SearchQuantumTasks requests per second (RPS) that you can send in one burst in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Burst rate of SearchSpendingLimits requests | Each supported Region: 50 | No | The maximum number of additional SearchSpendingLimits requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Burst rate of UpdateSpendingLimit requests | Each supported Region: 5 | No | The maximum number of additional UpdateSpendingLimit requests per second (RPS) that you can send in one burst in this account in the current Region. |
| Maximum allowed compute instances for a job | Each supported Region: 1 | Yes | The maximum allowed number of compute instances for a job. |
| Maximum number of instances of ml.c4.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.c4.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c4.4xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.c4.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.c4.8xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.c4.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c4.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.c4.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5.18xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.c5.18xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.c5.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5.4xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.c5.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.c5.9xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.c5.9xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.c5.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5n.18xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.c5n.18xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5n.2xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.c5n.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5n.4xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.c5n.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.c5n.9xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.c5n.9xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.c5n.xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.c5n.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g4dn.12xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.g4dn.12xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g4dn.16xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.g4dn.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g4dn.2xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.g4dn.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.g4dn.4xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.g4dn.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g4dn.8xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.g4dn.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g4dn.xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.g4dn.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g5.12xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.12xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g5.16xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.g5.24xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.24xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g5.2xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g5.48xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.48xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g5.4xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g5.8xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.g5.xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.g5.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.12xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6.12xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.16xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.24xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.g6.24xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.g6.48xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.g6.48xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.4xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.8xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6e.12xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.g6e.12xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of instances of ml.g6e.16xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.g6e.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6e.24xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.g6e.24xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6e.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6e.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6e.48xlarge for jobs | Each supported Region: 1 | Yes | The maximum number of instances of type ml.g6e.48xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6e.4xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6e.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.g6e.8xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6e.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.g6e.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.g6e.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m4.10xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.m4.10xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m4.16xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.m4.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m4.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.m4.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.m4.4xlarge for jobs | Each supported Region: 2 | Yes | The maximum number of instances of type ml.m4.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m4.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.m4.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m5.12xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.m5.12xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m5.24xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.m5.24xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m5.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.m5.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of instances of ml.m5.4xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.m5.4xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m5.large for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.m5.large allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.m5.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.m5.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.p2.16xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.p2.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.p2.8xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.p2.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.p2.xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.p2.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.p3.16xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.p3.16xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.p3.2xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.p3.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.p3.8xlarge for jobs | Each supported Region: 0 | Yes | The maximum number of instances of type ml.p3.8xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.p3dn.24xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.p3dn.24xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances of ml.p4d.24xlarge for jobs | Each supported Region: 0 | No | The maximum number of instances of type ml.p4d.24xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.t3.2xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.t3.2xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.t3.large for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.t3.large allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Maximum number of instances of ml.t3.xlarge for jobs | Each supported Region: 5 | Yes | The maximum number of instances of type ml.t3.xlarge allowed for all Amazon Braket Hybrid Jobs in this account and region. |
| Number of concurrent DM1 tasks | us-east-1: 100<br>us-west-2: 100<br>Each of the other supported Regions: 50 | No | The maximum number of concurrent tasks running on the Density Matrix Simulator (DM1) in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of concurrent SV1 tasks | us-east-1: 100<br><br>us-west-2: 100<br><br>Each of the other supported Regions: 50 | No | The maximum number of concurrent tasks running on the State Vector Simulator (SV1) in the current Region. |
| Number of concurrent TN1 tasks | eu-west-2: 5<br><br>Each of the other supported Regions: 10 | Yes | The maximum number of concurrent tasks running on the Tensor Network Simulator (TN1) in the current Region. |
| Number of concurrent jobs | Each supported Region: 3 | Yes | The maximum number of concurrent jobs running in the current Region. |
| Rate of API requests | Each supported Region: 140 | Yes | The maximum number of requests per second that you can send in this account in the current Region. |
| Rate of CancelJob requests | Each supported Region: 2 per second | Yes | The maximum number of CancelJob requests you can send per second in this account per Region. |
| Rate of CancelQuantumTask requests | Each supported Region: 2 per second | Yes | The maximum number of CancelQuantumTask requests you can send per second in this account per Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateJob requests | Each supported Region: 1 per second | Yes | The maximum number of CreateJob requests you can send per second in this account per Region. |
| Rate of CreateQuantumTask requests | Each supported Region: 20 per second | Yes | The maximum number of CreateQuantumTask requests you can send per second in this account per Region. |
| Rate of CreateSpendingLimit requests | Each supported Region: 2 per second | Yes | The maximum number of CreateSpendingLimit requests you can send per second in this account per Region. |
| Rate of DeleteSpendingLimit requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteSpendingLimit requests you can send per second in this account per Region. |
| Rate of GetDevice requests | Each supported Region: 5 per second | Yes | The maximum number of GetDevice requests you can send per second in this account per Region. |
| Rate of GetJob requests | Each supported Region: 5 per second | Yes | The maximum number of GetJob requests you can send per second in this account per Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetQuantumTask requests | Each supported Region: 100 per second | Yes | The maximum number of GetQuantumTask requests you can send per second in this account per Region. |
| Rate of SearchDevices requests | Each supported Region: 5 per second | Yes | The maximum number of SearchDevices requests you can send per second in this account per Region. |
| Rate of SearchJobs requests | Each supported Region: 5 per second | Yes | The maximum number of SearchJobs requests you can send per second in this account per Region. |
| Rate of SearchQuantumTasks requests | Each supported Region: 5 per second | Yes | The maximum number of SearchQuantumTasks requests you can send per second in this account per Region. |
| Rate of SearchSpendingLimits requests | Each supported Region: 5 per second | Yes | The maximum number of SearchSpendingLimits requests you can send per second in this account per Region. |
| Rate of UpdateSpendingLimit requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateSpendingLimit requests you can send per second in this account per Region. |

# AWS Certificate Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | acm.us-east-2.amazonaws.com | HTTPS |
| | | acm-fips.us-east-2.api.aws | HTTPS |
| | | acm.us-east-2.api.aws | HTTPS |
| | | acm-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | acm.us-east-1.amazonaws.com | HTTPS |
| | | acm-fips.us-east-1.api.aws | HTTPS |
| | | acm-fips.us-east-1.amazonaws.com | HTTPS |
| | | acm.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | acm.us-west-1.amazonaws.com | HTTPS |
| | | acm.us-west-1.api.aws | HTTPS |
| | | acm-fips.us-west-1.amazonaws.com | HTTPS |
| | | acm-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | acm.us-west-2.amazonaws.com | HTTPS |
| | | acm-fips.us-west-2.amazonaws.com | HTTPS |
| | | acm.us-west-2.api.aws | HTTPS |
| | | acm-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | acm.af-south-1.amazonaws.com | HTTPS |
| | | acm.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | acm.ap-east-1.amazonaws.com | HTTPS |
| | | acm.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | acm.ap-south-2.amazonaws.com | HTTPS |
| | | acm.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | acm.ap-southeast-3.amazonaws.com | HTTPS |
| | | acm.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | acm.ap-southeast-5.amazonaws.com | HTTPS |
| | | acm.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | acm.ap-southeast-4.amazonaws.com | HTTPS |
| | | acm.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | acm.ap-south-1.amazonaws.com<br>acm.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | acm.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | acm.ap-northeast-3.amazonaws.com<br>acm.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | acm.ap-northeast-2.amazonaws.com<br>acm.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | acm.ap-southeast-1.amazonaws.com<br>acm.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | acm.ap-southeast-2.amazonaws.com<br>acm.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | acm.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | acm.ap-southeast-7.amazonaws.com<br>acm.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | acm.ap-northeast-1.amazonaws.com | HTTPS |
| | | acm.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | acm.ca-central-1.amazonaws.com | HTTPS |
| | | acm.ca-central-1.api.aws | HTTPS |
| | | acm-fips.ca-central-1.amazonaws.com | HTTPS |
| | | acm-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | acm.ca-west-1.amazonaws.com | HTTPS |
| | | acm-fips.ca-west-1.api.aws | HTTPS |
| | | acm.ca-west-1.api.aws | HTTPS |
| | | acm-fips.ca-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | acm.eu-central-1.amazonaws.com | HTTPS |
| | | acm.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | acm.eu-west-1.amazonaws.com | HTTPS |
| | | acm.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | acm.eu-west-2.amazonaws.com | HTTPS |
| | | acm.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | acm.eu-south-1.amazonaws.com | HTTPS |
| | | acm.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | acm.eu-west-3.amazonaws.com | HTTPS |
| | | acm.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | acm.eu-south-2.amazonaws.com | HTTPS | |
| | | acm.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | acm.eu-north-1.amazonaws.com | HTTPS | |
| | | acm.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | acm.eu-central-2.amazonaws.com | HTTPS | |
| | | acm.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | acm.il-central-1.amazonaws.com | HTTPS | |
| | | acm.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | acm.mx-central-1.amazonaws.com | HTTPS | |
| | | acm.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | acm.me-south-1.amazonaws.com | HTTPS | |
| | | acm.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | acm.me-central-1.amazonaws.com | HTTPS | |
| | | acm.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | acm.sa-east-1.amazonaws.com | HTTPS | |
| | | acm.sa-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | acm.us-gov-east-1.amazonaws.com | HTTPS | |
| | | acm.us-gov-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | acm.us-gov-west-1.amazonaws.com<br><br>acm.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ACM certificates | Each supported Region: 2,500 | Yes | The maximum number of ACM Certificates you can have in this account in the current Region. Only certificates in the PENDING or ISSUED state count towards this limit. |
| ACM certificates created in last 365 days | Each supported Region: 5,000 | Yes | The maximum number of ACM Certificates you can request per year. |
| ACME domain validations per ACME endpoint | Each supported Region: 1,500 | Yes | The maximum number of ACME domain validations per ACME endpoint. |
| ACME endpoints per account | Each supported Region: 50 | Yes | The maximum number of ACME endpoints you can have in this account in the current Region. |
| Domain names per ACM certificate | Each supported Region: 10 | Yes | The maximum number of domain names per |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | ACM Certificate. The first domain name that you submit is included as the subject common name (CN) of the certificate. All names are included in the Subject Alternative Name extension. |
| External account bindings per ACME endpoint | Each supported Region: 1,000 | Yes | The maximum number of external account bindings (EABs) per ACME endpoint. |
| Identifiers per ACME order | Each supported Region: 100 | No | The maximum number of domain names (identifiers) per ACME order. |
| Imported certificates | Each supported Region: 2,500 | Yes | The maximum number of imported certificates you can have in this account in the current Region. |
| Imported certificates in last 365 days | Each supported Region: 5,000 | Yes | The maximum number of certificates you can import per year in this account in the current Region. |
| Rate of ChangeAccountKey requests (public CA) | Each supported Region: 50 per second | Yes | The maximum number of ChangeAccountKey requests per second to an ACME endpoint (Public CA). |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateAcmeDomainValidation API requests | Each supported Region: 50 per second | Yes | The maximum number of CreateAcmeDomainValidation API requests per second. |
| Rate of CreateAcmeEndpoint API requests | Each supported Region: 50 per second | Yes | The maximum number of CreateAcmeEndpoint API requests per second. |
| Rate of CreateAcmeExternalAccountBinding API requests | Each supported Region: 50 per second | Yes | The maximum number of CreateAcmeExternalAccountBinding API requests per second. |
| Rate of DeleteAcmeDomainValidation API requests | Each supported Region: 50 per second | Yes | The maximum number of DeleteAcmeDomainValidation API requests per second. |
| Rate of DeleteAcmeEndpoint API requests | Each supported Region: 50 per second | Yes | The maximum number of DeleteAcmeEndpoint API requests per second. |
| Rate of DeleteAcmeExternalAccountBinding API requests | Each supported Region: 50 per second | Yes | The maximum number of DeleteAcmeExternalAccountBinding API requests per second. |
| Rate of DescribeAcmeAccount API requests | Each supported Region: 50 per second | Yes | The maximum number of DescribeAcmeAccount API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DescribeAcmeDomainValidation API requests | Each supported Region: 50 per second | Yes | The maximum number of DescribeAcmeDomainValidation API requests per second. |
| Rate of DescribeAcmeEndpoint API requests | Each supported Region: 50 per second | Yes | The maximum number of DescribeAcmeEndpoint API requests per second. |
| Rate of DescribeAcmeExternalAccountBinding API requests | Each supported Region: 50 per second | Yes | The maximum number of DescribeAcmeExternalAccountBinding API requests per second. |
| Rate of FinalizeOrder requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of FinalizeOrder requests per second to an ACME endpoint (Public CA). |
| Rate of GetAcmeExternalAccountBindingCredentials API requests | Each supported Region: 50 per second | Yes | The maximum number of GetAcmeExternalAccountBindingCredentials API requests per second. |
| Rate of GetCertificate requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of GetCertificate requests per second to an ACME endpoint (Public CA). |
| Rate of GetDirectory requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of GetDirectory requests per second to an ACME endpoint (Public CA). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetOrder requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of GetOrder requests per second to an ACME endpoint (Public CA). |
| Rate of ListAcmeAccounts API requests | Each supported Region: 50 per second | Yes | The maximum number of ListAcmeAccounts API requests per second. |
| Rate of ListAcmeDomainValidations API requests | Each supported Region: 50 per second | Yes | The maximum number of ListAcmeDomainValidations API requests per second. |
| Rate of ListAcmeEndpoints API requests | Each supported Region: 50 per second | Yes | The maximum number of ListAcmeEndpoints API requests per second. |
| Rate of ListAcmeExternalAccountBindings API requests | Each supported Region: 50 per second | Yes | The maximum number of ListAcmeExternalAccountBindings API requests per second. |
| Rate of ListOrders requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of ListOrders requests per second to an ACME endpoint (Public CA). |
| Rate of ManageAccount requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of ManageAccount requests per second to an ACME endpoint (Public CA). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ManageAuthorization requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of ManageAuthorization requests per second to an ACME endpoint (Public CA). |
| Rate of NewAccount requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of NewAccount requests per second to an ACME endpoint (Public CA). |
| Rate of NewNonce requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of NewNonce requests per second to an ACME endpoint (Public CA). |
| Rate of NewOrder requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of NewOrder requests per second to an ACME endpoint (Public CA). |
| Rate of RestoreAcmeEndpoint API requests | Each supported Region: 50 per second | Yes | The maximum number of RestoreAcmeEndpoint API requests per second. |
| Rate of RevokeAcmeAccount API requests | Each supported Region: 50 per second | Yes | The maximum number of RevokeAcmeAccount API requests per second. |
| Rate of RevokeAcmeCertificate API requests | Each supported Region: 50 per second | Yes | The maximum number of RevokeAcmeCertificate API requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of RevokeAcmeExternalAccountBinding API requests | Each supported Region: 50 per second | Yes | The maximum number of RevokeAcmeExternalAccountBinding API requests per second. |
| Rate of RevokeCertificate requests (Public CA) | Each supported Region: 50 per second | Yes | The maximum number of RevokeCertificate requests per second to an ACME endpoint (Public CA). |
| Rate of UpdateAcmeDomainValidation API requests | Each supported Region: 50 per second | Yes | The maximum number of UpdateAcmeDomainValidation API requests per second. |
| Rate of UpdateAcmeEndpoint API requests | Each supported Region: 50 per second | Yes | The maximum number of UpdateAcmeEndpoint API requests per second. |

For more information, see Quotas in the *AWS Certificate Manager User Guide*.

## Amazon Chime endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Amazon Chime has a single endpoint that supports HTTPS: service.chime.aws.amazon.com.

## Service quotas

The following table lists additional quotas for Amazon Chime rooms and memberships.

| Resource | Default |
| --- | --- |
| Rooms per account | 1,500 |
| Rooms per profile | 1,500 |
| Memberships per room | 1,000 |
| Memberships per profile | 1,000 |

# Amazon Chime SDK endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Contents**


- Service endpoints
  - Identity
  - Media pipeline regional endpoints
  - Messaging
  - Voice
  - WebRTC media sessions

- Legacy
- Service quotas
  - Amazon Chime SDK call analytics quotas
  - Amazon Chime SDK Identity quotas
  - Amazon Chime SDK Media Pipeline quotas
  - Amazon Chime SDK Messaging quotas
  - SIP trunking and voice quotas
  - Amazon Chime SDK WebRTC quotas
  - Legacy quotas

# Service endpoints

## Identity

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | identity-chime.us-east-1.amazonaws.com | HTTPS |
| | | identity-chime.us-east-1.api.aws | HTTPS |
| | | identity-chime-fips.us-east-1.amazonaws.com | HTTPS |
| | | identity-chime-fips.us-east-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | identity-chime.eu-central-1.amazonaws.com | HTTPS |
| | | identity-chime.eu-central-1.api.aws | HTTPS |

# Media pipeline regional endpoints

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | media-pipelines-chime.us-east-1.amazonaws.com<br>media-pipelines-chime.us-east-1.api.aws<br>media-pipelines-chime-fips.us-east-1.amazonaws.com<br>media-pipelines-chime-fips.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | media-pipelines-chime.us-west-2.amazonaws.com<br>media-pipelines-chime.us-west-2.api.aws<br>media-pipelines-chime-fips.us-west-2.amazonaws.com<br>media-pipelines-chime-fips.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | media-pipelines-chime.ap-south-1.amazonaws.com<br>media-pipelines-chime.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | media-pipelines-chime.ap-northeast-2.amazonaws.com<br>media-pipelines-chime.ap-northeast-2.api.aws | HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | media-pipelines-chime.ap-southeast-1.amazonaws.com<br>media-pipelines-chime.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | media-pipelines-chime.ap-southeast-2.amazonaws.com<br>media-pipelines-chime.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | media-pipelines-chime.ap-northeast-1.amazonaws.com<br>media-pipelines-chime.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | media-pipelines-chime.ca-central-1.amazonaws.com<br>media-pipelines-chime.ca-central-1.api.aws<br>media-pipelines-chime-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | media-pipelines-chime.eu-central-1.amazonaws.com<br>media-pipelines-chime.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | media-pipelines-chime.eu-west-2.amazonaws.com<br>media-pipelines-chime.eu-west-2.api.aws | HTTPS<br>HTTPS |

## Messaging

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | messaging-chime.us-east-1.amazonaws.com<br>messaging-chime.us-east-1.api.aws<br>messaging-chime-fips.us-east-1.amazonaws.com<br>messaging-chime-fips.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | messaging-chime.eu-central-1.amazonaws.com<br>messaging-chime.eu-central-1.api.aws | HTTPS<br>HTTPS |

## Voice

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | voice-chime.us-east-1.amazonaws.com<br>voice-chime.us-east-1.api.aws<br>voice-chime-fips.us-east-1.amazonaws.com<br>voice-chime-fips.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | voice-chime.us-west-2.amazonaws.com<br>voice-chime.us-west-2.api.aws<br>voice-chime-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | voice-chime-fips.us-west-2.api.aws | |
| Asia Pacific (Seoul) | ap-northeast-2 | voice-chime.ap-northeast-2.amazonaws.com<br>voice-chime.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | voice-chime.ap-southeast-1.amazonaws.com<br>voice-chime.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | voice-chime.ap-southeast-2.amazonaws.com<br>voice-chime.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | voice-chime.ap-northeast-1.amazonaws.com<br>voice-chime.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | voice-chime.ca-central-1.amazonaws.com<br>voice-chime.ca-central-1.api.aws<br>voice-chime-fips.ca-central-1.amazonaws.com<br>voice-chime-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | voice-chime.eu-west-1.amazonaws.com<br>voice-chime.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | voice-chime.eu-central-1.amazonaws.com<br>voice-chime.eu-central-1.api.aws | HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (London) | eu-west-2 | voice-chime.eu-west-2.amazonaws.com<br>voice-chime.eu-west-2.api.aws | HTTPS<br>HTTPS |

## WebRTC media sessions

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | meetings-chime.us-east-1.amazonaws.com<br>meetings-chime.us-east-1.api.aws<br>meetings-chime-fips.us-east-1.amazonaws.com<br>meetings-chime-fips.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | meetings-chime.us-west-2.amazonaws.com<br>meetings-chime.us-west-2.api.aws<br>meetings-chime-fips.us-west-2.amazonaws.com<br>meetings-chime-fips.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | meetings-chime.af-south-1.amazonaws.com<br>meetings-chime.af-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | meetings-chime.ap-south-1.amazonaws.com<br>meetings-chime.ap-south-1.api.aws | HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Seoul) | `ap-northeast-2` | meetings-chime.ap-northeast-2.amazonaws.com<br>meetings-chime.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | `ap-southeast-1` | meetings-chime.ap-southeast-1.amazonaws.com<br>meetings-chime.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | `ap-southeast-2` | meetings-chime.ap-southeast-2.amazonaws.com<br>meetings-chime.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | `ap-northeast-1` | meetings-chime.ap-northeast-1.amazonaws.com<br>meetings-chime.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | `ca-central-1` | meetings-chime.ca-central-1.amazonaws.com<br>meetings-chime.ca-central-1.api.aws<br>meetings-chime-fips.ca-central-1.amazonaws.com<br>meetings-chime-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Frankfurt) | `eu-central-1` | meetings-chime.eu-central-1.amazonaws.com<br>meetings-chime.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | `eu-west-2` | meetings-chime.eu-west-2.amazonaws.com<br>meetings-chime.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | `il-central-1` | meetings-chime.il-central-1.amazonaws.com<br>meetings-chime.il-central-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | `us-gov-east-1` | meetings-chime.us-gov-east-1.amazonaws.com<br>meetings-chime.us-gov-east-1.api.aws<br>meetings-chime-fips.us-gov-east-1.amazonaws.com<br>meetings-chime-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | `us-gov-west-1` | meetings-chime.us-gov-west-1.amazonaws.com<br>meetings-chime.us-gov-west-1.api.aws<br>meetings-chime-fips.us-gov-west-1.amazonaws.com<br>meetings-chime-fips.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

### Legacy

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | service.chime.amazonaws.com | HTTPS |
| | | service-fips.chime.amazonaws.com | HTTPS |

## Service quotas

The tables in the following sections list and describe the quotas for the Amazon Chime SDK services.

> **Note**
>
> Service quotas are per AWS Region. If adjustable, quotas only change for the requested Region.

**Topics**


- Amazon Chime SDK call analytics quotas
- Amazon Chime SDK Identity quotas
- Amazon Chime SDK Media Pipeline quotas
- Amazon Chime SDK Messaging quotas
- SIP trunking and voice quotas
- Amazon Chime SDK WebRTC quotas
- Legacy quotas


### Amazon Chime SDK call analytics quotas

Amazon Chime SDK call analytics has the following quotas.

| Resource | Default limit | Adjustable | Description |
|---|---|---|---|
| Amazon Chime SDK call analytics - | 100 | Yes | Maximum number of call analytics |

| Resource | Default limit | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum configurations | | | configurations per Region for the specified account. |
| Amazon Chime SDK call analytics - Maximum pipelines | 20 | Yes | Maximum number of call analytics media processing pipelines per Region for the specified account. |
| Amazon Chime SDK call analytics - Maximum voice profile domains | 3 | Yes | Voice analytics voice profile domains per Region for the specified account. Customers can create voice profile domains to collect a set of related voice IDs for voice analytics. |
| Amazon Chime SDK call analytics - Voice profiles per voice profile domain | 20 | Yes | Maximum number of voice analytics voice profiles per voice profile domain. |
| Amazon Chime SDK call analytics - Maximum voice analytics tasks | 25 | Yes | Maximum number of active voice analytics tasks per Region to perform speaker search and voice tone analysis. |
| Amazon Chime SDK call analytics - Voice profile domain API rate | 0 | Yes | The maximum TPS for all voice profile domain API requests. |

| Resource | Default limit | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK call analytics - Voice profile API rate | 0 | Yes | The maximum TPS for all voice profile API requests. |
| Amazon Chime SDK call analytics - Speaker search API rate | 0 | Yes | The maximum TPS for all speaker search API requests. |
| Amazon Chime SDK call analytics - Voice tone analysis API rate | 0 | Yes | The maximum TPS for all voice tone analysis API requests. |
| Amazon Chime SDK call analytics - API rate | 5 | Yes | The maximum call analytics configura tion management API requests per second for this account in the current Region. |

**Note**

- If you exceed the quota for any Region, you receive a **Resource Limit Exceeded** exception. You can use the **Service Quotas** page in the AWS console to request an increase, or you can contact your customer support representative.
- Several of the call analytics APIs create resources and API requests for other AWS services. Those additional resources count against your account's quotas. If you request a quota or transactions-per-second increase from call analytics, you must also request increases for those other AWS services. Otherwise, your requests may be throttled and fail.

## Amazon Chime SDK Identity quotas

The following quotas apply for managing AppInstances, AppInstanceUsers and AppInstanceEndpoints for the Amazon Chime SDK Identity APIs. The quotas apply on a per API basis.

Amazon Chime SDK Identity has the following resource quotas:

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amazon Chime SDK Identity - Maximum `AppInstances` per AWS Account | 100 | Yes | The maximum number of `AppInstances` you can create in an AWS Account. |
| Amazon Chime SDK Identity - Maximum `AppInstanceUsers` per `AppInstance` | 100,000 | Yes | The maximum number of `AppInstanceUsers` you can create in an `AppInstance` |
| Amazon Chime SDK Identity - Maximum `AppInstanceUserAdmins` per `AppInstance` | 100 | Yes | The maximum number of `AppInstanceUserAdmins` you can create in an `AppInstance` . |
| Amazon Chime SDK Identity - Maximum AppInstanceUserEndpoints per AppInstanceUser | 10 | Yes | The maximum number of `AppInstanceUserEndpoints` you can create for an `AppInstanceUser` |

The following limits apply to the Amazon Chime SDK Identity APIs for managing `AppInstances`, `AppInstanceUsers`, and `AppInstanceEndpoints` on a per-API basis.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK Identity - Rate of CreateAppInstance, UpdateAppInstance, and DeleteAppInstance requests | 5 requests per second per `AppInstance` | Yes | The maximum requests per second at which you can simultaneously create, update and delete `AppInstances` in your AWS account |
| Amazon Chime SDK Identity - Rate of CreateAppInstanceUser, UpdateAppInstanceUser, and DeleteAppInstanceUser requests | 10 requests per second per `AppInstance` | Yes | The maximum requests per second at which you can simultaneously create, update and delete `AppInstanceUsers` in your AWS account |
| Amazon Chime SDK Identity - Rate of DescribeAppInstance requests | 20 requests per second per AppInstance | Yes | The maximum requests per second at which you can invoke the DescribeAppInstance API for a specific `AppInstance` |
| Amazon Chime SDK Identity - Rate of DescribeAppInstanceUser requests | 20 requests per second per `AppInstanceUser` | Yes | The maximum requests per second at which you can invoke the DescribeAppInstanceUser |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | API for a specific `AppInstanceUser` |
| Amazon Chime SDK Identity - Rate of ListAppInstances requests | 10 requests per second per AWS Account | Yes | The maximum requests per second at which you can invoke the ListAppInstances API for your AWS Account |
| Amazon Chime SDK Identity - Rate of ListAppInstanceUsers requests | 10 requests per second per AppInstanceUser | Yes | The maximum requests per second at which you can invoke the ListAppInstanceUsers API for a specific `AppInstance` |
| Amazon Chime SDK Identity - Total requests per second per AWS Account | 50 requests per second per AWS account | Yes | The maximum requests per second at which you can simultaneously invoke the Chime SDK Identity APIs in your AWS account |

## Amazon Chime SDK Media Pipeline quotas

Amazon Chime SDK Media Pipelines have the following quotas.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK media pipeline - API rate | 10 | Yes | The maximum rate of all SDK media pipeline API requests |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | per second that you can send in this account in the current Region. |
| Amazon Chime SDK media pipeline - Maximum pipelines | `US East (N. Virginia)` : 100<br><br>All others: 10 | Yes | The maximum number of concurrent media pipelines that you can run for this account. Excludes media insights pipelines. |
| Amazon Chime SDK media pipeline - Maximum Amazon Kinesis Video Stream pools | 10 | Yes | The maximum number of Amazon Kinesis Video Stream (KVS) pools that you can create. A pool contains multiple KVS streams. The total number of KVS streams that can be created is governed by the KVS service quota: Number of video streams. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK media pipeline - Amazon Kinesis Video Stream pool API rate | 5 | Yes | The maximum number of requests per second that you make to each of the following APIs: CreateMediaPipelineKinesisVideoStreamPool, UpdateMediaPipelineKinesisVideoStreamPool, and DeleteMediaPipelineKinesisVideoStreamPool. |

## Amazon Chime SDK Messaging quotas

Amazon Chime SDK Messaging has the following quotas.

| Name | Default | Adjustable |
| --- | --- | --- |
| Amazon Chime SDK Messaging - Maximum concurrent active connections per AppInstanceUser | 10 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelFlows per AppInstance | 100 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelMemberships per Channel | 10,000 | Yes |

| Name | Default | Adjustable |
| --- | --- | --- |
| Amazon Chime SDK Messaging - Maximum ChannelProcessors per ChannelFlow | 1 | Yes |
| Amazon Chime SDK Messaging - Maximum Channels per AppInstance | 10,000,000 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelModerators per Channel | 1,000 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelProcessors per ChannelFlow | 1 | Yes |
| Amazon Chime SDK Messaging - Maximum CHANNEL_DETAILS events for prefetch | 50 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelMessages in CHANNEL_DETAILS events for prefetch | 20 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelMemberships in CHANNEL_DETAILS events for prefetch | 30 | Yes |

| Name | Default | Adjustable |
| --- | --- | --- |
| Amazon Chime SDK Messaging - Maximum Elastic Channels per AppInstance * | 1 | Yes |
| Amazon Chime SDK Messaging - Maximum SubChannels per Elastic Channel * | 100 | Yes |
| Amazon Chime SDK Messaging - Maximum ChannelMemberships per SubChannel * | 1,000 | Yes<br>**Note**<br>The product of `MaximumSubChannelsPerElastiCChannel` and `MaximumChannelMembershipsPerSubChannel` cannot exceed 1-million after a limit increase. |
| Amazon Chime SDK Messaging – Maximum ChannelMemberships per AppInstanceUser | 20,000 | Yes |

*Only available in the US East (N. Virginia) Region.

**Note**

Customers with large chat channels often qualify for volume discounts on pricing. Contact your account manager for more details.

In addition to the quotas listed above, Amazon Chime SDK Messaging has the following API rates.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK Messaging - Rate of CreateChannel, UpdateChannel, DeleteChannel requests | 15 requests per second per `AppInstance` | Yes | The maximum requests per second at which you can simultaneously create, update, or delete Channels for a specific `AppInstance` |
| Amazon Chime SDK Messaging - Rate of ListChannels requests | 10 requests per second per `AppInstance` | Yes | The maximum rate at which you can invoke the ListChannels API for a specific `AppInstance` |
| Amazon Chime SDK Messaging - Rate of DescribeChannel requests | 20 requests per second per Channel | Yes | The maximum rate at which you can invoke the DescribeChannel API for a specific `AppInstance` |
| Amazon Chime SDK Messaging - Rate of `CreateChannelMembership` , `DeleteChannelMembership` , `CreateCha` | `us-east-1` : 15 requests per second per Channel, 30 requests per second per Elastic Channel | Yes | The maximum requests per second at which you can simultaneously create, update, or delete `Channel` memberships, bans, |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| nnelBan , DeleteCha nnelBan , CreateCha nnelModer ator , DeleteCha nnelModerator requests | eu-central-1 : 10 requests per second per Channel | | and moderators for a specific channel |
| Amazon Chime SDK Messaging - Rate of ListChann elMemberships , ListChannelBans , ListChann elModerators requests | 15 requests per second per Channel and Elastic Channel | Yes | The maximum requests per second at which you can simultaneously list Channel membershi ps, bans, and moderators for a specific channel |
| Amazon Chime SDK Messaging - Rate of SendChann elMessage , UpdateCha nnelMessa ge , RedactCha nnelMessa ge , DeleteCha nnelMessage requests | us-east-1 : 30 requests per second per Channel or per SubChannel .<br><br>eu-central-1 : 10 requests per second per channel | Yes | The maximum requests per second at which you can simultaneously create, update, redact, or delete ChannelMe ssages for a specific Channel or SubChannel . This includes calls to the ChannelFl owCallback API to update messages. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK Messaging - Rate of ListChannelMessages requests | 30 requests per second per Channel or per SubChannel | Yes | The maximum requests per second at which you can invoke ListChannelMessages for a specific Channel or SubChannel . |
| Amazon Chime SDK Messaging - Rate of GetChannelMessage requests | 30 requests per second per Channel or per sub-channel | Yes | The maximum requests per second at which you can invoke GetChannelMessage for a specific Channel or SubChannel . |
| Amazon Chime SDK Messaging - Rate of ListChannelMembershipsForAppInstanceUser , ListChannelsModeratedByAppInstanceUser requests | 15 requests per second per AppInstanceUser | Yes | The maximum requests per second at which you can simultaneously invoke the ListChannelMembershipsForAppInstanceUser and ListChannelsModeratedByAppInstanceUser APIs for a specific AppInstance |
| Amazon Chime SDK Messaging - Rate of ListSubChannels requests | 15 requests per second per elastic channel | Yes | The maximum requests per second at which you can invoke the ListSubChannels API for a specific Elastic Channel |

Amazon Chime SDK Messaging also has the following general API rates.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK Messaging - Rate of API requests per ChimeBearer | us-east-1 : 10 requests per second per AppInstanceUser<br><br>eu-central-1 : 5 requests per second per AppInstanceUser | Yes | The maximum requests per second at which a specific ChimeBearer can simultaneously invoke the Messaging APIs. This limit prevents a single AppInstanceUser from consuming all of the throughput for operations on a single Channel in an AppInstance . You can request limit increases for specific AppInstanceUsers . |
| Amazon Chime SDK Messaging - Rate of Batch API requests | 1 request per second per resource | No | The maximum requests per second at which each Batch API, such as BatchCreateChannelMembership, can be invoked |
| Amazon Chime SDK Messaging - Rate of websocket connect requests per AppInstanceUser | 1 request per second per AppInstanceUser | No | The maximum requests per second at which an AppInstanceUser can establish a new websocket connectio |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | n, even for different sessions |

## SIP trunking and voice quotas

Amazon Chime SDK SIP trunking and voice have the following quotas.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK SIP trunking and voice - provisioned phone numbers | 5 | Yes | The maximum number of provisioned phone numbers you can have in this account |
| Amazon Chime SDK SIP trunking and voice - Voice Connectors | 3 | Yes | The maximum number of Voice Connectors you can have in the current AWS Region |
| Amazon Chime SDK SIP trunking and voice - Voice Connector groups | 3 | Yes | The maximum number of Voice Connector groups you can have in this account. |
| Amazon Chime SDK SIP trunking and voice - Voice Connectors per Voice Connector group | 3 | Yes | The maximum number of Voice Connectors you can have in any Voice Connector group |
| Amazon Chime SDK SIP trunking and | 1 | Yes | The maximum calls per second (CPS) allowed on each |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| voice - calls per second (CPS) | | | Voice Connector in the current AWS Region |
| Amazon Chime SDK SIP trunking and voice - SIP media applications | 30 | Yes | The maximum number of SIP media applications you can have in the current AWS Region |
| Amazon Chime SDK SIP trunking and voice - SIP media applications per SIP rule | 25 | Yes | The maximum number of SIP media applications you can have in a SIP rule in the current AWS Region |
| Amazon Chime SDK SIP trunking and voice - CreateSip MediaApplicationCall API rate limit | 1 | Yes | The maximum number of requests to the CreateSip MediaApplicationCall API per second allowed in the current AWS Region. |
| Amazon Chime SDK SIP trunking and voice - UpdateSip MediaApplicationCall API rate limit | 5 | Yes | The maximum number of requests to the UpdateSip MediaApplicationCall API per second allowed in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK SIP trunking and voice - StartBotConversation Amazon Lex bots | 0 | Yes | The maximum number of Amazon Lex bots you can use with the PSTN Audio StartBotConversation action in the current AWS Region. |
| Amazon Chime SDK SIP trunking and voice - Voice Connector Active Call Count | 10 | Yes | The maximum Voice Connector concurrent call count allowed per account in the current AWS Region. |
| Amazon Chime SDK SIP trunking and voice - SIP Media Application Active Call Count | 10 | Yes | The maximum PSTN Audio concurrent call count allowed per account in the current AWS Region. |

## Amazon Chime SDK WebRTC quotas

Amazon Chime SDK Meetings WebRTC media sessions have the following quotas.

| Name | Default | Adjustable |
| --- | --- | --- |
| Amazon Chime SDK Meetings – Concurrent meeting limit | 250 | Yes |
| Amazon Chime SDK meetings - Attendees per meeting | 250 | No |
| Amazon Chime SDK meetings - Replica meetings per primary meeting | 4 | Yes For more information, refer to the Chime SDK Meetings - replica meetings |

| Name | Default | Adjustable |
| --- | --- | --- |
| | | per primary meeting page in the Service Quotas console. |
| Amazon Chime SDK meetings - Concurrent video streams published per meeting | 25 | Yes. For more informati on, refer to the Chime SDK Meetings - maximum concurrent video streams published per meeting page in the Service Quotas console. |
| Amazon Chime SDK meetings - Concurrent video streams subscribed per attendee | 25 | No |
| Amazon Chime SDK meetings - BatchCreateAttendees API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |
| Amazon Chime SDK meetings - CreateAttendee API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |
| Amazon Chime SDK meetings - CreateMeeting API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |
| Amazon Chime SDK meetings - CreateMeetingWithA ttendees API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |
| Amazon Chime SDK meetings - DeleteAttendee API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |

| Name | Default | Adjustable |
| --- | --- | --- |
| Amazon Chime SDK meetings - DeleteMeeting API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |
| Amazon Chime SDK meetings - GetMeeting API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |
| Amazon Chime SDK meetings - ListAttendees API rate in requests per second | 10 | Yes. Adjusts automatically when concurrent meeting limit is adjusted. |

## Legacy quotas

The following quotas apply only to legacy endpoints.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK media pipeline - Maximum pipelines | 10 | No | The maximum rate of all media pipeline API requests that you can send in this account in the current Region. This quota is for the Amazon Chime namespace. |
| Amazon Chime – Total active media pipelines per account | 100 | No | The total number of active media pipelines per account in a Region. This quota is for the Amazon Chime namespace. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Amazon Chime SDK Media Pipelines – Rate limit for al media pipeline API requests in transacti ons per second | 10 | Yes | The maximum rate of all SDK media pipeline API requests that you can send in this account in the current Region This quota is for the Amazon Chime namespace. |
| Amazon Chime SDK Media Pipeline - Media capture API burst limit | 10 | No | The maximum number of media capture pipeline requests that you can send in one burst. This quota is for the Amazon Chime namespace . This quota is for the Amazon Chime namespace. |
| Amazon Chime SDK Media Pipeline - API rate limit | 10 | Yes | The maximum rate of all SDK media pipeline API requests that you can send in this account in the current Region. This quota is for the Amazon Chime namespace. |

# AWS Clean Rooms endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | cleanrooms.us-east-2.amazonaws.com | HTTPS |
| | | cleanrooms-fips.us-east-2.api.aws | HTTPS |
| | | cleanrooms.us-east-2.api.aws | HTTPS |
| | | cleanrooms-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | cleanrooms.us-east-1.amazonaws.com | HTTPS |
| | | cleanrooms-fips.us-east-1.api.aws | HTTPS |
| | | cleanrooms-fips.us-east-1.amazonaws.com | HTTPS |
| | | cleanrooms.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | cleanrooms.us-west-2.amazonaws.com | HTTPS |
| | | cleanrooms-fips.us-west-2.api.aws | HTTPS |
| | | cleanrooms.us-west-2.api.aws | HTTPS |
| | | cleanrooms-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | cleanrooms.ap-northeast-2.amazonaws.com<br>cleanrooms.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | cleanrooms.ap-southeast-1.amazonaws.com<br>cleanrooms.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | cleanrooms.ap-southeast-2.amazonaws.com<br>cleanrooms.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | cleanrooms.ap-northeast-1.amazonaws.com<br>cleanrooms.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | cleanrooms.eu-central-1.amazonaws.com<br>cleanrooms.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | cleanrooms.eu-west-1.amazonaws.com<br>cleanrooms.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | cleanrooms.eu-west-2.amazonaws.com<br>cleanrooms.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | cleanrooms.eu-north-1.amazonaws.com<br>cleanrooms.eu-north-1.api.aws | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Analysis rule size | Each supported Region: 100 Kilobytes | No | Maximum size of JSON for an analysis rule |
| Analysis templates per membership | Each supported Region: 100 | Yes | Maximum number of analysis templates per membership |
| CR.1X workers per PySpark job | Each supported Region: 128 | Yes | Maximum number of CR.1X workers per PySpark job |
| CR.4X workers per PySpark job | Each supported Region: 32 | Yes | Maximum number of CR.4X workers per PySpark job |
| Collaborations created per account | Each supported Region: 500 | Yes | Maximum number of collaborations created per account |
| Columns per configured table allowlist | Each supported Region: 100 | Yes | Maximum number of columns that can be allowlisted per configured table |
| Concurrent PySpark job vCPU usage per account | us-east-1: 1,024<br><br>Each of the other supported Regions: 512 | Yes | Maximum total vCPU usage of all concurrently running PySpark jobs per account |
| Concurrent PySpark jobs per account | us-east-1: 5 | Yes | Maximum number of concurrently running PySpark jobs per account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | Each of the other supported Regions: 2 | | |
| Concurrent SQL queries per account | us-east-1: 5<br><br>Each of the other supported Regions: 2 | Yes | Maximum number of concurrently running SQL queries per account |
| Concurrent SQL query vCPU usage per account | Each supported Region: 512 | Yes | Maximum total vCPU usage of all concurrently running SQL queries per account |
| Concurrent ongoing job per membership | Each supported Region: 1 | No | Maximum number of concurrent ongoing jobs per membership |
| Concurrent ongoing log exports per account | Each supported Region: 10 | Yes | Maximum number of concurrent ongoing log exports per account |
| Concurrent ongoing log exports per membership | Each supported Region: 5 | Yes | Maximum number of concurrent ongoing log exports per membership |
| Concurrent ongoing queries per membership | Each supported Region: 5 | No | Maximum number of concurrent ongoing queries per membership |
| Configured audience model associations per membership | Each supported Region: 5 | Yes | Maximum number of configured audience model associations per membership |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Configured tables per account | Each supported Region: 250 | Yes | Maximum number of configured tables created per account |
| Configured tables per protected query | Each supported Region: 15 | Yes | Maximum number of configured tables in a protected query |
| ID mapping tables per membership | Each supported Region: 5 | Yes | Maximum number of ID mapping tables per membership |
| ID namespace associations per membership | Each supported Region: 10 | Yes | Maximum number of ID namespace associations per membership |
| Intermediate tables per membership | Each supported Region: 25 | Yes | Maximum number of intermediate tables per membership |
| Members invited per collaboration | Each supported Region: 5 | Yes | Maximum number of members invited per collaboration |
| Memberships per account | Each supported Region: 500 | Yes | Maximum number of memberships per account |
| Privacy budget templates for access budgets per collaboration | Each supported Region: 25 | Yes | Maximum number of privacy budget templates for access budgets per collaboration |
| Query text length (using differential privacy) | Each supported Region: 8 Kilobytes | No | Maximum text length for a SQL query statement using differential privacy |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Query text length on Clean Rooms SQL analytics engine | Each supported Region: 90 Kilobytes | No | Maximum text length for a SQL query statement on Clean Rooms SQL analytics engine |
| Query text length on Spark analytics engine | Each supported Region: 500 Kilobytes | No | Maximum text length for a SQL query statement on Spark analytics engine |
| Rate of BatchGetCollaborationAnalysisTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of BatchGetCollaborationAnalysisTemplate API calls per second |
| Rate of BatchGetSchema requests | Each supported Region: 5 per second | Yes | Maximum number of BatchGetSchema API calls per second |
| Rate of BatchGetSchemaAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of BatchGetSchemaAnalysisRule API calls per second |
| Rate of CreateAnalysisTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of CreateAnalysisTemplate API calls per second |
| Rate of CreateCollaboration requests | Each supported Region: 5 per second | Yes | Maximum number of CreateCollaboration API calls per second |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateCollaborationChangeRequest requests | Each supported Region: 5 per second | Yes | Maximum number of CreateCollaborationChangeRequest API calls per second |
| Rate of CreateConfiguredAudienceModelAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of CreateConfiguredAudienceModelAssociation API calls per second |
| Rate of CreateConfiguredTable requests | Each supported Region: 5 per second | Yes | Maximum number of CreateConfiguredTable API calls per second |
| Rate of CreateConfiguredTableAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of CreateConfiguredTableAnalysisRule API calls per second |
| Rate of CreateConfiguredTableAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of CreateConfiguredTableAssociation API calls per second |
| Rate of CreateConfiguredTableAssociationAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of CreateConfiguredTableAssociationAnalysisRule API calls per second |
| Rate of CreateIdMappingTable requests | Each supported Region: 5 per second | Yes | Maximum number of CreateIdMappingTable API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateIdNamespaceAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of CreateIdNamespaceAssociation API calls per second |
| Rate of CreateMembership requests | Each supported Region: 5 per second | Yes | Maximum number of CreateMembership API calls per second |
| Rate of CreatePrivacyBudgetTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of CreatePrivacyBudgetTemplate API calls per second |
| Rate of DeleteAnalysisTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteAnalysisTemplate API calls per second |
| Rate of DeleteCollaboration requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteCollaboration API calls per second |
| Rate of DeleteConfiguredAudienceModelAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteConfiguredAudienceModelAssociation API calls per second |
| Rate of DeleteConfiguredTable requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteConfiguredTable API calls per second |
| Rate of DeleteConfiguredTableAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteConfiguredTableAnalysisRule API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DeleteConfiguredTableAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteConfiguredTableAssociation API calls per second |
| Rate of DeleteConfiguredTableAssociationAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteConfiguredTableAssociationAnalysisRule API calls per second |
| Rate of DeleteIdMappingTable requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteIdMappingTable API calls per second |
| Rate of DeleteIdNamespaceAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteIdNamespaceAssociation API calls per second |
| Rate of DeleteMember requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteMember API calls per second |
| Rate of DeleteMembership requests | Each supported Region: 5 per second | Yes | Maximum number of DeleteMembership API calls per second |
| Rate of DeletePrivacyBudgetTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of DeletePrivacyBudgetTemplate API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetAnalysisLogExport requests | Each supported Region: 5 per second | Yes | Maximum number of GetAnalysisLogExport API calls per second |
| Rate of GetAnalysisTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of GetAnalysisTemplate API calls per second |
| Rate of GetCollaboration requests | Each supported Region: 5 per second | Yes | Maximum number of GetCollaboration API calls per second |
| Rate of GetCollaborationAnalysisTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of GetCollaborationAnalysisTemplate API calls per second |
| Rate of GetCollaborationChangeRequest requests | Each supported Region: 5 per second | Yes | Maximum number of GetCollaborationChangeRequest API calls per second |
| Rate of GetCollaborationConfiguredAudienceModelAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of GetCollaborationConfiguredAudienceModelAssociation API calls per second |
| Rate of GetCollaborationIdNamespaceAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of GetCollaborationIdNamespaceAssociation API calls per second |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetCollaborationPrivacyBudgetTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of GetCollaborationPrivacyBudgetTemplate API calls per second |
| Rate of GetConfiguredAudienceModelAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of GetConfiguredAudienceModelAssociation API calls per second |
| Rate of GetConfiguredTable requests | Each supported Region: 20 per second | Yes | Maximum number of GetConfiguredTable API calls per second |
| Rate of GetConfiguredTableAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of GetConfiguredTableAnalysisRule API calls per second |
| Rate of GetConfiguredTableAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of GetConfiguredTableAssociation API calls per second |
| Rate of GetConfiguredTableAssociationAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of GetConfiguredTableAssociationAnalysisRule API calls per second |
| Rate of GetIdMappingTable requests | Each supported Region: 5 per second | Yes | Maximum number of GetIdMappingTable API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetIdNamespaceAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of GetIdNamespaceAssociation API calls per second |
| Rate of GetMembership requests | Each supported Region: 5 per second | Yes | Maximum number of GetMembership API calls per second |
| Rate of GetPrivacyBudgetTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of GetPrivacyBudgetTemplate API calls per second |
| Rate of GetProtectedJob requests | Each supported Region: 5 per second | Yes | Maximum number of GetProtectedJob API calls per second |
| Rate of GetProtectedQuery requests | Each supported Region: 20 per second | Yes | Maximum number of GetProtectedQuery API calls per second |
| Rate of GetSchema requests | Each supported Region: 5 per second | Yes | Maximum number of GetSchema API calls per second |
| Rate of GetSchemaAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of GetSchemaAnalysisRule API calls per second |
| Rate of ListAnalysisLogExports requests | Each supported Region: 5 per second | Yes | Maximum number of ListAnalysisLogExports API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListAnalysisTemplates requests | Each supported Region: 5 per second | Yes | Maximum number of ListAnalysisTemplates API calls per second |
| Rate of ListCollaborationAnalysisTemplates requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborationAnalysisTemplates API calls per second |
| Rate of ListCollaborationChangeRequests requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborationChangeRequests API calls per second |
| Rate of ListCollaborationConfiguredAudienceModelAssociations requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborationConfiguredAudienceModelAssociations API calls per second |
| Rate of ListCollaborationIdNamespaceAssociations requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborationIdNamespaceAssociations API calls per second |
| Rate of ListCollaborationPrivacyBudgetTemplates requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborationPrivacyBudgetTemplates API calls per second |
| Rate of ListCollaborationPrivacyBudgets requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborationPrivacyBudgets API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListCollaborations requests | Each supported Region: 5 per second | Yes | Maximum number of ListCollaborations API calls per second |
| Rate of ListConfiguredAudienceModelAssociations requests | Each supported Region: 5 per second | Yes | Maximum number of ListConfiguredAudienceModelAssociations API calls per second |
| Rate of ListConfiguredTableAssociations requests | Each supported Region: 5 per second | Yes | Maximum number of ListConfiguredTableAssociations API calls per second |
| Rate of ListConfiguredTables requests | Each supported Region: 5 per second | Yes | Maximum number of ListConfiguredTables API calls per second |
| Rate of ListIdMappingTables requests | Each supported Region: 5 per second | Yes | Maximum number of ListIdMappingTables API calls per second |
| Rate of ListIdNamespaceAssociations requests | Each supported Region: 5 per second | Yes | Maximum number of ListIdNamespaceAssociations API calls per second |
| Rate of ListMembers requests | Each supported Region: 5 per second | Yes | Maximum number of ListMembers API calls per second |
| Rate of ListMemberships requests | Each supported Region: 5 per second | Yes | Maximum number of ListMemberships API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListPrivacyBudgetTemplates requests | Each supported Region: 5 per second | Yes | Maximum number of ListPrivacyBudgetTemplates API calls per second |
| Rate of ListPrivacyBudgets requests | Each supported Region: 5 per second | Yes | Maximum number of ListPrivacyBudgets API calls per second |
| Rate of ListProtectedJobs requests | Each supported Region: 5 per second | Yes | Maximum number of ListProtectedJobs API calls per second |
| Rate of ListProtectedQueries requests | Each supported Region: 5 per second | Yes | Maximum number of ListProtectedQueries API calls per second |
| Rate of ListSchemas requests | Each supported Region: 5 | Yes | Maximum number of ListSchemas API calls per second |
| Rate of ListTagsForResource requests | Each supported Region: 5 per second | Yes | Maximum number of ListTagsForResource API calls per second |
| Rate of PopulateIdMappingTable requests | Each supported Region: 5 per second | Yes | Maximum number of PopulateIdMappingTable API calls per second |
| Rate of PreviewPrivacyImpact requests | Each supported Region: 5 per second | Yes | Maximum number of PreviewPrivacyImpact API calls per second |
| Rate of StartAnalysisLogExport requests | Each supported Region: 5 per second | Yes | Maximum number of StartAnalysisLogExport API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of StartProtectedJob requests | Each supported Region: 5 per second | Yes | Maximum number of StartProtectedJob API calls per second |
| Rate of StartProtectedQuery requests | Each supported Region: 5 per second | Yes | Maximum number of StartProtectedQuery API calls per second |
| Rate of TagResource requests | Each supported Region: 5 per second | Yes | Maximum number of TagResource API calls per second |
| Rate of UntagResource requests | Each supported Region: 5 per second | Yes | Maximum number of UntagResource API calls per second |
| Rate of UpdateAnalysisTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateAnalysisTemplate API calls per second |
| Rate of UpdateCollaboration requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateCollaboration API calls per second |
| Rate of UpdateCollaborationChangeRequest requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateCollaborationChangeRequest API calls per second |
| Rate of UpdateConfiguredAudienceModelAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateConfiguredAudienceModelAssociation API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of UpdateConfiguredTable requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateConfiguredTable API calls per second |
| Rate of UpdateConfiguredTableAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateConfiguredTableAnalysisRule API calls per second |
| Rate of UpdateConfiguredTableAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateConfiguredTableAssociation API calls per second |
| Rate of UpdateConfiguredTableAssociationAnalysisRule requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateConfiguredTableAssociationAnalysisRule API calls per second |
| Rate of UpdateIdMappingTable requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateIdMappingTable API calls per second |
| Rate of UpdateIdNamespaceAssociation requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateIdNamespaceAssociation API calls per second |
| Rate of UpdatePrivacyBudgetTemplate requests | Each supported Region: 5 per second | Yes | Maximum number of UpdatePrivacyBudgetTemplate API calls per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of UpdateProtectedJob requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateProtectedJob API calls per second |
| Rate of UpdateProtectedQuery requests | Each supported Region: 5 per second | Yes | Maximum number of UpdateProtectedQuery API calls per second |
| Table associations per membership | Each supported Region: 100 | Yes | Maximum number of table associations per membership |

For more information, see AWS Clean Rooms quotas in the *AWS Clean Rooms User Guide*.

# AWS Clean Rooms ML endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cleanrooms-ml.us-east-2.amazonaws.com<br><br>cleanrooms-ml.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | cleanrooms-ml.us-east-1.amazonaws.com<br>cleanrooms-ml.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | cleanrooms-ml.us-west-2.amazonaws.com<br>cleanrooms-ml.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cleanrooms-ml.ap-northeast-2.amazonaws.com<br>cleanrooms-ml.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cleanrooms-ml.ap-southeast-1.amazonaws.com<br>cleanrooms-ml.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cleanrooms-ml.ap-southeast-2.amazonaws.com<br>cleanrooms-ml.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cleanrooms-ml.ap-northeast-1.amazonaws.com<br>cleanrooms-ml.ap-northeast-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cleanrooms-ml.eu-central-1.amazonaws.com<br>cleanrooms-ml.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | cleanrooms-ml.eu-west-1.amazonaws.com<br>cleanrooms-ml.eu-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | cleanrooms-ml.eu-west-2.amazonaws.com<br><br>cleanrooms-ml.eu-west-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | cleanrooms-ml.eu-north-1.amazonaws.com<br><br>cleanrooms-ml.eu-north-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Active audience export jobs per audience generation job | Each supported Region: 25 | No | The maximum number of active audience export jobs for an audience generation job |
| Active configured model algorithm associations per membership | Each supported Region: 1,000 | Yes | The maximum number of active configured model algorithm associations per membership |
| Active configured model algorithms per membership | Each supported Region: 1,000 | Yes | The maximum number of active configured model algorithms per membership |
| Active custom model input channels per membership | Each supported Region: 100 | Yes | The maximum number of active custom model input channels per membership |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of active training instances per account | Each supported Region: 10 | Yes | The maximum number of active training instances that can be used per account for creating trained models. |
| Maximum number of active/pending/in progress trained model versions per trained model | Each supported Region: 100 | Yes | The maximum number of active/pending/in progress trained model versions that can be created per trained model. |
| Maximum number of input columns to generate synthetic data | Each supported Region: 1,000 | No | The maximum number of input columns to generate synthetic data |
| Maximum number of input rows to generate synthetic data | Each supported Region: 2,500,000 | No | The maximum number of input rows to generate synthetic data |
| Maximum number of ml.c4.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c4.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c4.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c4.4xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.c4.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c4.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c4.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c4.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5.18xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5.18xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5.9xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5.9xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of ml.c5.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5n.18xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5n.18xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5n.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5n.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5n.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5n.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5n.9xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5n.9xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c5n.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c5n.xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of ml.c6i.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c6i.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c6i.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c6i.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c6i.32xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.32xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c6i.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.4xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.c6i.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c6i.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c6i.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.2xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of ml.c7i.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.48xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.large training instances that can be used per account for creating trained models. |
| Maximum number of ml.c7i.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.c7i.xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of ml.g4dn.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g4dn.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g4dn.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g4dn.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g4dn.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g4dn.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g4dn.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g4dn.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g4dn.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g4dn.8xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.g4dn.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g4dn.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.48xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.g5.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g5.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g5.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.24xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.g6.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.48xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6e.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.12xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.g6e.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6e.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6e.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6e.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.48xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6e.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.g6e.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.8xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.g6e.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.g6e.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m4.10xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m4.10xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m4.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m4.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m4.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m4.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m4.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m4.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m4.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m4.xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.m5.12xlarge training instances per account | Each supported Region: 3 | Yes | The maximum number of ml.m5.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m5.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m5.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m5.2xlarge training instances per account | Each supported Region: 3 | Yes | The maximum number of ml.m5.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m5.4xlarge training instances per account | Each supported Region: 3 | Yes | The maximum number of ml.m5.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m5.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m5.large training instances that can be used per account for creating trained models. |
| Maximum number of ml.m5.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m5.xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of ml.m6i.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.32xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.32xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.4xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.m6i.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.large training instances that can be used per account for creating trained models. |
| Maximum number of ml.m6i.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m6i.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.24xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.m7i.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.48xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.large training instances that can be used per account for creating trained models. |
| Maximum number of ml.m7i.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.m7i.xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.p2.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p2.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.p2.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p2.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.p2.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p2.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.p4d.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p4d.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.p4de.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p4de.24xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of ml.p5.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p5.48xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.p5en.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.p5en.48xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.24xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.r5.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.large training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5d.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.12xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.r5d.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5d.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5d.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5d.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5d.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r5d.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.large training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.r5d.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r5d.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.12xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.12xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.16xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.16xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.24xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.24xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.48xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.48xlarge training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.r7i.4xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.4xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.8xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.8xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.large training instances that can be used per account for creating trained models. |
| Maximum number of ml.r7i.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.r7i.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.t3.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.t3.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.t3.large training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.t3.large training instances that can be used per account for creating trained models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ml.t3.medium training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.t3.medium training instances that can be used per account for creating trained models. |
| Maximum number of ml.t3.xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.t3.xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.trn1.2xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.trn1.2xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.trn1.32xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.trn1.32xlarge training instances that can be used per account for creating trained models. |
| Maximum number of ml.trn1n.32xlarge training instances per account | Each supported Region: 0 | Yes | The maximum number of ml.trn1n.32xlarge training instances that can be used per account for creating trained models. |
| Maximum unique values in categorical prediction column | Each supported Region: 100 | No | The maximum unique values in categorical prediction column |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum number of input columns required to generate synthetic data | Each supported Region: 5 | No | The minimum number of input columns required to generate synthetic data |
| Minimum number of input rows required to generate synthetic data | Each supported Region: 1,500 | No | The minimum number of input rows required to generate synthetic data |
| Pending/In-progress audience export jobs per customer | Each supported Region: 20 | No | The maximum number of pending/in-progress audience export jobs per customer |
| Pending/In-progress audience generation jobs per customer | Each supported Region: 10 | Yes | The maximum number of pending/in-progress audience generation jobs per customer |
| Pending/In-progress audience models per customer | Each supported Region: 2 | Yes | The maximum number of pending/in-progress audience model training jobs per customer |
| Pending/In-progress custom model inference jobs per account | Each supported Region: 10 | Yes | The maximum number of pending/in-progress custom model inference jobs per account |
| Pending/In-progress custom model inference jobs per membership | Each supported Region: 5 | Yes | The maximum number of pending/in-progress custom model inference jobs per membership |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Pending/In-progress custom model training jobs per account | Each supported Region: 10 | Yes | The maximum number of pending/in-progress custom model training jobs per account |
| Pending/In-progress custom model training jobs per membership | Each supported Region: 5 | Yes | The maximum number of pending/in-progress custom model training jobs per membership |
| Pending/In-progress synthetic MLIC generation jobs per account | Each supported Region: 2 | Yes | The maximum number of pending/in-progress synthetic MLIC generation jobs per account |

For more information, see AWS Clean Rooms ML quotas in the *AWS Clean Rooms ML User Guide*.

## Cloud Control API endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cloudcontrolapi.us-east-2.amazonaws.com | HTTPS | |
| | | cloudcontrolapi-fips.us-east-2.api.aws | HTTPS | |
| | | cloudcontrolapi-fips.us-east-2.amazonaws.com | HTTPS | |
| | | cloudcontrolapi.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | cloudcontrolapi.us-east-1.amazonaws.com | HTTPS | |
| | | cloudcontrolapi-fips.us-east-1.api.aws | HTTPS | |
| | | cloudcontrolapi-fips.us-east-1.amazonaws.com | HTTPS | |
| | | cloudcontrolapi.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | cloudcontrolapi.us-west-1.amazonaws.com | HTTPS | |
| | | cloudcontrolapi-fips.us-west-1.api.aws | HTTPS | |
| | | cloudcontrolapi-fips.us-west-1.amazonaws.com | HTTPS | |
| | | cloudcontrolapi.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | cloudcontrolapi.us-west-2.amazonaws.com | HTTPS | |
| | | cloudcontrolapi-fips.us-west-2.api.aws | HTTPS | |
| | | cloudcontrolapi-fips.us-west-2.amazonaws.com | HTTPS | |
| | | cloudcontrolapi.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | cloudcontrolapi.af-south-1.amazonaws.com<br>cloudcontrolapi.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cloudcontrolapi.ap-east-1.amazonaws.com<br>cloudcontrolapi.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | cloudcontrolapi.ap-south-2.amazonaws.com<br>cloudcontrolapi.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | cloudcontrolapi.ap-southeast-3.amazonaws.com<br>cloudcontrolapi.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | cloudcontrolapi.ap-southeast-5.amazonaws.com<br>cloudcontrolapi.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | cloudcontrolapi.ap-southeast-4.amazonaws.com<br>cloudcontrolapi.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cloudcontrolapi.ap-south-1.amazonaws.com<br>cloudcontrolapi.ap-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | cloudcontrolapi.ap-southeast-6.amazonaws.com<br>cloudcontrolapi.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | cloudcontrolapi.ap-northeast-3.amazonaws.com<br>cloudcontrolapi.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | cloudcontrolapi.ap-northeast-2.amazonaws.com<br>cloudcontrolapi.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudcontrolapi.ap-southeast-1.amazonaws.com<br>cloudcontrolapi.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudcontrolapi.ap-southeast-2.amazonaws.com<br>cloudcontrolapi.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | cloudcontrolapi.ap-east-2.amazonaws.com<br>cloudcontrolapi.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | cloudcontrolapi.ap-southeast-7.amazonaws.com<br>cloudcontrolapi.ap-southeast-7.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudcontrolapi.ap-northeast-1.amazonaws.com<br>cloudcontrolapi.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | cloudcontrolapi.ca-central-1.amazonaws.com<br>cloudcontrolapi-fips.ca-central-1.api.aws<br>cloudcontrolapi-fips.ca-central-1.amazonaws.com<br>cloudcontrolapi.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | cloudcontrolapi.ca-west-1.amazonaws.com<br>cloudcontrolapi-fips.ca-west-1.api.aws<br>cloudcontrolapi-fips.ca-west-1.amazonaws.com<br>cloudcontrolapi.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | cloudcontrolapi.eu-central-1.amazonaws.com<br>cloudcontrolapi.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | cloudcontrolapi.eu-west-1.amazonaws.com<br>cloudcontrolapi.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | cloudcontrolapi.eu-west-2.amazonaws.com<br>cloudcontrolapi.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | cloudcontrolapi.eu-south-1.amazonaws.com<br>cloudcontrolapi.eu-south-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | cloudcontrolapi.eu-west-3.amazonaws.com<br>cloudcontrolapi.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | cloudcontrolapi.eu-south-2.amazonaws.com<br>cloudcontrolapi.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | cloudcontrolapi.eu-north-1.amazonaws.com<br>cloudcontrolapi.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | cloudcontrolapi.eu-central-2.amazonaws.com<br>cloudcontrolapi.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | cloudcontrolapi.il-central-1.amazonaws.com<br>cloudcontrolapi.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | cloudcontrolapi.mx-central-1.amazonaws.com<br>cloudcontrolapi.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | cloudcontrolapi.me-south-1.amazonaws.com<br>cloudcontrolapi.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | cloudcontrolapi.me-central-1.amazonaws.com<br>cloudcontrolapi.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | cloudcontrolapi.sa-east-1.amazonaws.com<br>cloudcontrolapi.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | cloudcontrolapi.us-gov-east-1.amazonaws.com<br>cloudcontrolapi-fips.us-gov-east-1.api.aws<br>cloudcontrolapi-fips.us-gov-east-1.amazonaws.com<br>cloudcontrolapi.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | cloudcontrolapi.us-gov-west-1.amazonaws.com<br>cloudcontrolapi-fips.us-gov-west-1.api.aws<br>cloudcontrolapi-fips.us-gov-west-1.amazonaws.com<br>cloudcontrolapi.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

### Service quotas

This service has no quotas.

## AWS Cloud9 endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cloud9.us-east-2.amazonaws.com | HTTPS | |
| | | cloud9-fips.us-east-2.api.aws | HTTPS | |
| | | cloud9-fips.us-east-2.amazonaws.com | HTTPS | |
| | | cloud9.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | cloud9.us-east-1.amazonaws.com | HTTPS | |
| | | cloud9-fips.us-east-1.api.aws | HTTPS | |
| | | cloud9-fips.us-east-1.amazonaws.com | HTTPS | |
| | | cloud9.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | cloud9.us-west-1.amazonaws.com | HTTPS | |
| | | cloud9-fips.us-west-1.api.aws | HTTPS | |
| | | cloud9-fips.us-west-1.amazonaws.com | HTTPS | |
| | | cloud9.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | cloud9.us-west-2.amazonaws.com | HTTPS | |
| | | cloud9-fips.us-west-2.api.aws | HTTPS | |
| | | cloud9-fips.us-west-2.amazonaws.com | HTTPS | |
| | | cloud9.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | cloud9.af-south-1.amazonaws.com | HTTPS | |
| | | cloud9.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | cloud9.ap-east-1.amazonaws.com<br>cloud9.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cloud9.ap-south-1.amazonaws.com<br>cloud9.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cloud9.ap-northeast-3.amazonaws.com<br>cloud9.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cloud9.ap-northeast-2.amazonaws.com<br>cloud9.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cloud9.ap-southeast-1.amazonaws.com<br>cloud9.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cloud9.ap-southeast-2.amazonaws.com<br>cloud9.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cloud9.ap-northeast-1.amazonaws.com<br>cloud9.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | cloud9.ca-central-1.amazonaws.com | HTTPS |
| | | cloud9-fips.ca-central-1.api.aws | HTTPS |
| | | cloud9-fips.ca-central-1.amazonaws.com | HTTPS |
| | | cloud9.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | cloud9.eu-central-1.amazonaws.com | HTTPS |
| | | cloud9.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | cloud9.eu-west-1.amazonaws.com | HTTPS |
| | | cloud9.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | cloud9.eu-west-2.amazonaws.com | HTTPS |
| | | cloud9.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | cloud9.eu-south-1.amazonaws.com | HTTPS |
| | | cloud9.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | cloud9.eu-west-3.amazonaws.com | HTTPS |
| | | cloud9.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | cloud9.eu-north-1.amazonaws.com | HTTPS |
| | | cloud9.eu-north-1.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | cloud9.il-central-1.amazonaws.com | HTTPS |
| | | cloud9.il-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | cloud9.me-south-1.amazonaws.com<br>cloud9.me-south-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | cloud9.sa-east-1.amazonaws.com<br>cloud9.sa-east-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| EC2 development environments for this account | Each supported Region: 200 | Yes | The maximum number of AWS Cloud9 EC2 development environments that you can create in this account in the current Region. |
| EC2 development environments for this user | Each supported Region: 100 | Yes | The maximum number of AWS Cloud9 EC2 development environments that you can create for this user in the current Region. |
| Members per development environment | Each supported Region: 8 | No | The maximum number of users that can participate in a single AWS Cloud9 development environme |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | nt in this account in the current Region. |
| SSH development environments for this account | Each supported Region: 200 | Yes | The maximum number of AWS Cloud9 SSH development environments that you can create in this account in the current Region. |
| SSH development environments for this user | Each supported Region: 100 | Yes | The maximum number of AWS Cloud9 SSH development environments that you can create for this user in the current Region. |

For more information, see Quotas in the *AWS Cloud9 User Guide*.

## Amazon Cloud Directory endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | clouddirectory.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | clouddirectory.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | clouddirectory.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | clouddirectory.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | clouddirectory.ap-southeast-2.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | clouddirectory.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | clouddirectory.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | clouddirectory.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | clouddirectory.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud | us-gov-west-1 | clouddirectory.us-gov-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (US-West) | | | | |

## Service quotas

For more information, see Amazon Cloud Directory quotas.

# AWS Cloud WAN endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

AWS Cloud WAN is available in the following Regions. Cloud WAN aggregates and stores information in its home Region. Cloud WAN supports the following home Region: `us-west-2` — US West (Oregon).

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (London) | eu-west-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | networkmanager.us-gov-west-1.amazonaws.com<br>networkmanager.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | networkmanager.us-gov-west-1.amazonaws.com<br>networkmanager.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

For a list of quotas, see Quotas in the *AWS Cloud WAN User Guide*.

# AWS CloudFormation endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cloudformation.us-east-2.amazonaws.com | HTTPS | |
| | | cloudformation-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | cloudformation.us-east-1.amazonaws.com | HTTPS | |
| | | cloudformation-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | cloudformation.us-west-1.amazonaws.com | HTTPS | |
| | | cloudformation-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | cloudformation.us-west-2.amazonaws.com | HTTPS | |
| | | cloudformation-fips.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | cloudformation.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific | ap-east-1 | cloudformation.ap-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | cloudformation.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | cloudformation.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | cloudformation.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | cloudformation.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cloudformation.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | cloudformation.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cloudformation.ap-northeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | cloudformation.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudformation.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudformation.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | cloudformation.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | cloudformation.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudformation.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | cloudformation.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | cloudformation.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cloudformation.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | cloudformation.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cloudformation.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | cloudformation.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | cloudformation.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | cloudformation.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | cloudformation.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | cloudformation.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | cloudformation.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | cloudformation.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | cloudformation.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | cloudformation.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | cloudformation.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | cloudformation.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | cloudformation.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Custom Resource Response Data | Each supported Region: 4,096 Bytes | No | The maximum size of response data that a custom resource provider can submit in a response |
| Hook configuration size | Each supported Region: 204.8 Kilobytes | No | The maximum amount of configuration data a Hook can store. |
| Hooks per account | Each supported Region: 100 | No | The maximum number of Hooks allowed for this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Hooks per resource | Each supported Region: 100 | No | The maximum number of Hooks that can be configured per resource. |
| Maximum dependencies per stack set | Each supported Region: 10 per 5 minutes | No | Maximum number of auto-deployment dependencies allowed per stack set. |
| Maximum stack set dependencies per account | Each supported Region: 100 | No | Maximum number of auto-deployment dependencies allowed across all stack sets in your account. |
| Module limit per account | Each supported Region: 100 | Yes | Maximum number of module types you are allowed to register. |
| Nested modules | Each supported Region: 3 | No | Number of levels of nesting per module. |
| Queued operations per stack set | Each supported Region: 10,000 | No | Maximum number of queued operations for a stack set at a given time. |
| Resource limit per account | Each supported Region: 50 | Yes | Maximum number of resource types you are allowed to register. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Stack instance operations per administrator account | Each supported Region: 10,000 | Yes | Maximum number of stack instances, across all stack sets, that you can run operations on in each Region at the same time, per administrator account. |
| Stack instances per stack set | Each supported Region: 100,000 | Yes | Maximum number of stack instances you can create per stack set. |
| Stack sets per administrator account | Each supported Region: 1,000 | Yes | Maximum number of AWS CloudFormation stack sets you can create in your administrator account. |
| Stacks | Each supported Region: 2,000 | Yes | The maximum number of active stacks; an active stack is any stack with a status other than DELETE_COMPLETE |
| Stacks imported using S3 object per stack set operation | Each supported Region: 200 | No | Maximum number of stacks you can import using S3 object per stack set operation. |
| Stacks imported using inline stack ids per stack set operation | Each supported Region: 10 | No | Maximum number of stacks you can import using inline stack ids per stack set operation. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Template Description Length | Each supported Region: 1,024 Bytes | No | The maximum length of a CloudFormation template description |
| Template Dynamic References | Each supported Region: 60 | No | The maximum number of dynamic references per CloudFormation template |
| Template Logical Resource ID Length | Each supported Region: 255 | No | The maximum length of a CloudFormation template logical resource ID |
| Template Mapping Attributes | Each supported Region: 200 | No | The maximum number of attributes per CloudFormation template mapping |
| Template Mapping Name and Mapping Attribute Name Length | Each supported Region: 255 | No | The maximum length of a CloudFormation template mapping name and mapping attribute name |
| Template Mappings | Each supported Region: 200 | No | The maximum number of mappings per CloudFormation template |
| Template Output Name Length | Each supported Region: 255 | No | The maximum length of a CloudFormation template output name |
| Template Outputs | Each supported Region: 200 | No | The maximum number of outputs per CloudFormation template |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Template Parameter Name Length | Each supported Region: 255 | No | The maximum length of a CloudFormation template parameter name |
| Template Parameter Value Length | Each supported Region: 4,096 | No | The maximum length of a CloudFormation template parameter value |
| Template Parameters | Each supported Region: 200 | No | The maximum number of parameters per CloudFormation template |
| Template Resources | Each supported Region: 500 | No | The maximum number of resources per CloudFormation template |
| Template Size | Each supported Region: 1 Megabytes | No | The maximum size of a CloudFormation template; a template of size greater than 50 KiB must be stored in an Amazon S3 bucket and submitted in requests as an Amazon S3 object URL |
| Throttle rate limit for ActivateOrganizationsAccess | Each supported Region: 2 per second | No | The maximum number of ActivateOrganizationsAccess requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for BatchDescribeTypeConfigurations | Each supported Region: 2 per second | No | The maximum number of BatchDescribeTypeConfigurations requests you can make per second per account, in the current region |
| Throttle rate limit for CancelUpdateStack | Each supported Region: 2 per second | No | The maximum number of CancelUpdateStack requests you can make per second per account, in the current region |
| Throttle rate limit for ContinueUpdateRollback | Each supported Region: 2 per second | No | The maximum number of ContinueUpdateRollback requests you can make per second per account, in the current region |
| Throttle rate limit for CreateChangeSet | Each supported Region: 2 per second | No | The maximum number of CreateChangeSet requests you can make per second per account, in the current region |
| Throttle rate limit for CreateGeneratedTemplate | Each supported Region: 2 per second | No | The maximum number of CreateGeneratedTemplate requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for CreateStack | Each supported Region: 2 per second | No | The maximum number of CreateStack requests you can make per second per account, in the current region |
| Throttle rate limit for CreateStackInstances | Each supported Region: 2 per second | No | The maximum number of CreateStackInstances requests you can make per second per account, in the current region |
| Throttle rate limit for CreateStackSet | Each supported Region: 2 per second | No | The maximum number of CreateStackSet requests you can make per second per account, in the current region |
| Throttle rate limit for DeleteChangeSet | Each supported Region: 2 per second | No | The maximum number of DeleteChangeSet requests you can make per second per account, in the current region |
| Throttle rate limit for DeleteGeneratedTemplate | Each supported Region: 2 per second | No | The maximum number of DeleteGeneratedTemplate requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for DeleteStack | Each supported Region: 2 per second | No | The maximum number of DeleteStack requests you can make per second per account, in the current region |
| Throttle rate limit for DeleteStackInstances | Each supported Region: 2 per second | No | The maximum number of DeleteStackInstances requests you can make per second per account, in the current region |
| Throttle rate limit for DeleteStackSet | Each supported Region: 2 per second | No | The maximum number of DeleteStackSet requests you can make per second per account, in the current region |
| Throttle rate limit for DeregisterType | Each supported Region: 2 per second | No | The maximum number of DeregisterType requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeAccountLimits | Each supported Region: 2 per second | No | The maximum number of DescribeAccountLimits requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeChangeSet | Each supported Region: 2 per second | No | The maximum number of DescribeChangeSet requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for DescribeGeneratedTemplate | Each supported Region: 2 per second | No | The maximum number of DescribeGeneratedTemplate requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeOrganizationsAccess | Each supported Region: 2 per second | No | The maximum number of DescribeOrganizationsAccess requests you can make per second per account, in the current region |
| Throttle rate limit for DescribePublisher | Each supported Region: 2 per second | No | The maximum number of DescribePublisher requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeResourceScan | Each supported Region: 2 per second | No | The maximum number of DescribeResourceScan requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStackDriftDetectionStatus | Each supported Region: 2 per second | No | The maximum number of DescribeStackDriftDetectionStatus requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for DescribeStackEvents | Each supported Region: 10 per second | No | The maximum number of DescribeStackEvents requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStackInstance | Each supported Region: 2 per second | No | The maximum number of DescribeStackInstance requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStackResource | Each supported Region: 20 per second | No | The maximum number of DescribeStackResource requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStackResourceDrifts | Each supported Region: 2 per second | No | The maximum number of DescribeStackResourceDrifts requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStackResources | Each supported Region: 10 per second | No | The maximum number of DescribeStackResources requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for DescribeStackSet | Each supported Region: 2 per second | No | The maximum number of DescribeStackSet requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStackSetOperation | Each supported Region: 2 per second | No | The maximum number of DescribeStackSetOperation requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeStacks | Each supported Region: 10 per second | No | The maximum number of DescribeStacks requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeType | Each supported Region: 2 per second | No | The maximum number of DescribeType requests you can make per second per account, in the current region |
| Throttle rate limit for DescribeTypeRegistration | Each supported Region: 2 per second | No | The maximum number of DescribeTypeRegistration requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for DetectStackDrift | Each supported Region: 2 per second | No | The maximum number of DetectStackDrift requests you can make per second per account, in the current region |
| Throttle rate limit for DetectStackResourceDrift | Each supported Region: 3 per second | No | The maximum number of DetectStackResourceDrift requests you can make per second per account, in the current region |
| Throttle rate limit for DetectStackSetDrift | Each supported Region: 2 per second | No | The maximum number of DetectStackSetDrift requests you can make per second per account, in the current region |
| Throttle rate limit for EstimateTemplateCost | Each supported Region: 2 per second | No | The maximum number of EstimateTemplateCost requests you can make per second per account, in the current region |
| Throttle rate limit for ExecuteChangeSet | Each supported Region: 2 per second | No | The maximum number of ExecuteChangeSet requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for GetGeneratedTemplate | Each supported Region: 2 per second | No | The maximum number of GetGeneratedTemplate requests you can make per second per account, in the current region |
| Throttle rate limit for GetStackPolicy | Each supported Region: 2 per second | No | The maximum number of GetStackPolicy requests you can make per second per account, in the current region |
| Throttle rate limit for GetTemplate | Each supported Region: 10 per second | No | The maximum number of GetTemplate requests you can make per second per account, in the current region |
| Throttle rate limit for GetTemplateSummary | Each supported Region: 2 per second | No | The maximum number of GetTemplateSummary requests you can make per second per account, in the current region |
| Throttle rate limit for ListChangeSets | Each supported Region: 2 per second | No | The maximum number of ListChangeSets requests you can make per second per account, in the current region |
| Throttle rate limit for ListExports | Each supported Region: 10 per second | No | The maximum number of ListExports requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for ListGeneratedTemplates | Each supported Region: 2 per second | No | The maximum number of ListGeneratedTemplates requests you can make per second per account, in the current region |
| Throttle rate limit for ListImports | Each supported Region: 2 per second | No | The maximum number of ListImports requests you can make per second per account, in the current region |
| Throttle rate limit for ListResourceScanRelatedResources | Each supported Region: 2 per second | No | The maximum number of ListResourceScanRelatedResources requests you can make per second per account, in the current region |
| Throttle rate limit for ListResourceScanResources | Each supported Region: 2 per second | No | The maximum number of ListResourceScanResources requests you can make per second per account, in the current region |
| Throttle rate limit for ListResourceScans | Each supported Region: 2 per second | No | The maximum number of ListResourceScans requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for ListStackInstanceResourceDrifts | Each supported Region: 2 per second | No | The maximum number of ListStackInstanceResourceDrifts requests you can make per second per account, in the current region |
| Throttle rate limit for ListStackInstances | Each supported Region: 2 per second | No | The maximum number of ListStackInstances requests you can make per second per account, in the current region |
| Throttle rate limit for ListStackResources | Each supported Region: 10 per second | No | The maximum number of ListStackResources requests you can make per second per account, in the current region |
| Throttle rate limit for ListStackSetAutoDeploymentTargets | Each supported Region: 2 per second | No | The maximum number of ListStackSetAutoDeploymentTargets requests you can make per second per account, in the current region |
| Throttle rate limit for ListStackSetOperationResults | Each supported Region: 2 per second | No | The maximum number of ListStackSetOperationResults requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for ListStackSetOperations | Each supported Region: 2 per second | No | The maximum number of ListStackSetOperations requests you can make per second per account, in the current region |
| Throttle rate limit for ListStackSets | Each supported Region: 2 per second | No | The maximum number of ListStackSets requests you can make per second per account, in the current region |
| Throttle rate limit for ListStacks | Each supported Region: 10 per second | No | The maximum number of ListStacks requests you can make per second per account, in the current region |
| Throttle rate limit for ListTypeRegistrations | Each supported Region: 2 per second | No | The maximum number of ListTypeRegistrations requests you can make per second per account, in the current region |
| Throttle rate limit for ListTypeVersions | Each supported Region: 2 per second | No | The maximum number of ListTypeVersions requests you can make per second per account, in the current region |
| Throttle rate limit for ListTypes | Each supported Region: 2 per second | No | The maximum number of ListTypes requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for PublishType | Each supported Region: 2 per second | No | The maximum number of PublishType requests you can make per second per account, in the current region |
| Throttle rate limit for RegisterType | Each supported Region: 0.5 per second | No | The maximum number of RegisterType requests you can make per second per account, in the current region |
| Throttle rate limit for RollbackStack | Each supported Region: 2 per second | No | The maximum number of RollbackStack requests you can make per second per account, in the current region |
| Throttle rate limit for SetStackPolicy | Each supported Region: 2 per second | No | The maximum number of SetStackPolicy requests you can make per second per account, in the current region |
| Throttle rate limit for SetTypeDefaultVersion | Each supported Region: 2 per second | No | The maximum number of SetTypeDefaultVersion requests you can make per second per account, in the current region |
| Throttle rate limit for SignalResource | Each supported Region: 2 per second | No | The maximum number of SignalResource requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate limit for StartResourceScan | Each supported Region: 2 per second | No | The maximum number of StartResourceScan requests you can make per second per account, in the current region |
| Throttle rate limit for StopStackSetOperation | Each supported Region: 2 per second | No | The maximum number of StopStackSetOperation requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateGeneratedTemplate | Each supported Region: 2 per second | No | The maximum number of UpdateGeneratedTemplate requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateStack | Each supported Region: 2 per second | No | The maximum number of UpdateStack requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateStackInstances | Each supported Region: 2 per second | No | The maximum number of UpdateStackInstances requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for UpdateStackSet | Each supported Region: 2 per second | No | The maximum number of UpdateStackSet requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateTerminationProtection | Each supported Region: 2 per second | No | The maximum number of UpdateTerminationProtection requests you can make per second per account, in the current region |
| Throttle rate limit for ValidateTemplate | Each supported Region: 2 per second | No | The maximum number of ValidateTemplate requests you can make per second per account, in the current region |
| Version limit per module | Each supported Region: 100 | Yes | Maximum number of versions per module type you are allowed to register. |
| Version limit per resource | Each supported Region: 50 | Yes | Maximum number of versions per resource type you are allowed to register. |
| Versions per hook | Each supported Region: 100 | No | The maximum number of versions you can create per Hook. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Wait Condition Signal Data Size | Each supported Region: 4,096 Bytes | No | The maximum size of signal data that cfn-signal can submit in a wait condition signal |

For more information, see CloudFormation Quotas in the *AWS CloudFormation User Guide*.

# Amazon CloudFront endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Amazon CloudFront

| Region Name | Region | Endpoint | Protocol | Amazon Route 53 Hosted Zone ID* |
| --- | --- | --- | --- | --- |
| US East (N. Virginia) Region | us-east-1 | cloudfront.amazonaws.com<br>cloudfront-fips.amazonaws.com | HTTPS<br>HTTPS | Z2FDTNDATAQYW2 |
| | | cloudfront.global.api.aws | HTTPS | Z045326210N9XYH9VV43R |

| Region Name | Region | Endpoint | Protocol | Amazon Route 53 Hosted Zone ID* |
|---|---|---|---|---|
| | | cloudfront-fips.global.api.aws | HTTPS | Z0169911SZ2QDT2H60C8 |

### Amazon CloudFront KeyValueStore

CloudFront KeyValueStore is a global service. Specify your AWS account ID for the endpoint.

| Endpoint | Protocol | Amazon Route 53 Hosted Zone ID* |
|---|---|---|
| *AWS account*.cloudfront-kvs.global.api.aws | HTTPS | The root host zone is the following. Each customer endpoint has its own hosted zone.<br>Z024645223JLTKL21A7KJ |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Alternate domain names (CNAMEs) per distribution | Each supported Region: 100 | Yes | The maximum number of alternate domain names (CNAMEs) per distribution. |
| Anycast static IP lists per AWS account | Each supported Region: 0 | Yes | The maximum number of Anycast static IP lists per AWS account. |
| Cache behaviors per distribution | Each supported Region: 75 | Yes | The maximum number of cache behaviors per distribution. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Cache policies per AWS account | Each supported Region: 20 | Yes | The maximum number of cache policies per AWS account. |
| CloudFront Functions: Maximum number of distributions associated with a single function | Each supported Region: 100 | Yes | The maximum number of CloudFront distributions associated with a single CloudFront function. |
| Connection attempts per origin | Each supported Region: 3 | No | The number of connection attempts per origin (1-3). |
| Connection timeout per origin | Each supported Region: 10 Seconds | No | The connection timeout per origin (1-10 seconds). |
| Continuous deployment policies per AWS account | Each supported Region: 20 | Yes | The maximum number of continuous deployment policies that you can have for an AWS account. |
| Cookies per cache policy | Each supported Region: 10 | Yes | The maximum number of cookies per cache policy. |
| Cookies per origin request policy | Each supported Region: 10 | Yes | The maximum number of cookies per origin request policy. |
| Custom headers per response headers policy | Each supported Region: 10 | Yes | The maximum number of custom headers per response headers policy. |
| Custom headers: maximum length of a header name | Each supported Region: 256 | No | The maximum length of a header name in characters. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom headers: maximum length of a header value | Each supported Region: 1,783 | No | The maximum length of a header value in characters. |
| Custom headers: maximum length of all header values and names combined | Each supported Region: 10,240 | No | The maximum length of all header values and names combined. |
| Custom headers: maximum number of custom headers that you can configure CloudFront to add to origin requests | Each supported Region: 30 | Yes | The maximum number of custom headers that you can configure CloudFront to add to origin requests. |
| Data transfer rate per distribution | Each supported Region: 150 | Yes | The maximum data transfer rate (in Gbps) per distribution. |
| Distributions associated with a single key group | Each supported Region: 100 | Yes | The maximum number of distributions associated with a single key group. |
| Distributions associated with the same VPC origin | Each supported Region: 50 | No | The maximum number of distributions that can be associated with a single VPC origin. |
| Distributions associated with the same cache policy | Each supported Region: 100 | No | The maximum number of distributions associated with the same cache policy. |
| Distributions associated with the same origin access control | Each supported Region: 100 | Yes | The maximum number of distributions that you can associate with an origin access control. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Distributions associated with the same origin request policy | Each supported Region: 100 | No | The maximum number of distributions associated with the same origin request policy. |
| Distributions associated with the same response headers policy | Each supported Region: 100 | Yes | The maximum number of distributions that you can associate with the same response headers policy. |
| Distributions per AWS account that you can create triggers for | Each supported Region: 500 | Yes | The maximum number of distributions per AWS account that you can create triggers for. |
| File invalidation: maximum number of active wildcard invalidations allowed | Each supported Region: 15 | No | The maximum number of active wildcard invalidations allowed. |
| File invalidation: maximum number of files allowed in active invalidation requests, excluding wildcard invalidations | Each supported Region: 3,000 | No | The maximum number of files allowed in active invalidation requests, excluding wildcard invalidations. |
| Function memory size (Viewer request and response event) | Each supported Region: 128 Megabytes | No | The maximum function memory size (in MB). (Viewer request and response event) |
| Function timeout (Origin request and response event) | Each supported Region: 30 Seconds | No | The maximum function timeout (in seconds). (Origin request and response event) |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Function timeout for a viewer request and response event | Each supported Region: 5 Seconds | No | The maximum function timeout (in seconds). (Viewer request and response event) |
| Headers per cache policy | Each supported Region: 10 | Yes | The maximum number of headers per cache policy. |
| Headers per origin request policy | Each supported Region: 10 | Yes | The maximum number of headers per origin request policy. |
| Keep-alive timeout for a custom origin associated with distribution | Each supported Region: 120 Seconds | Yes | How long (in seconds) CloudFront tries to maintain a connection to your distributions custom origin after it gets the last packet of a response. |
| Key groups associated with a single distribution | Each supported Region: 4 | Yes | The maximum number of key groups associated with a single distribution. |
| Key groups per AWS account | Each supported Region: 10 | Yes | The maximum number of key groups per AWS account. |
| Length of the content security policy value in a response headers policy | Each supported Region: 1,783 | Yes | The maximum length of the content security policy value in a response headers policy. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum CA bundle size per trust store | Each supported Region: 64 Kilobytes | Yes | The maximum size (in KB) of a CA bundle that can be uploaded to a trust store. |
| Maximum certificate size in trust store CA bundle | Each supported Region: 16 Kilobytes | Yes | The maximum size (in KB) of an individual certificate within a trust stores CA bundle. |
| Maximum compressed size of a Lambda function and any included libraries. (Origin request and response event) | Each supported Region: 50 Megabytes | No | The maximum compressed size (in MB) of a Lambda function and any included libraries. (Origin request and response event) |
| Maximum compressed size of a Lambda function and any included libraries. (Viewer request and response event) | Each supported Region: 50 Megabytes | No | The maximum compressed size (in MB) of a Lambda function and any included libraries. (Viewer request and response event) |
| Maximum file size for HTTP GET, POST, and PUT requests | Each supported Region: 50 Gigabytes | No | The maximum file size (in GB) for HTTP GET, POST, and PUT requests. |
| Maximum length of a URL | Each supported Region: 8,192 Bytes | No | The maximum length of a URL (in bytes). |
| Maximum length of a field to encrypt | Each supported Region: 16 Kilobytes | No | The maximum length (in KB) of a field to encrypt. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum length of a request body when field-level encryption is configured | Each supported Region: 1 Megabytes | No | The maximum length (in MB) of a request body when field-level encryption is configured. |
| Maximum length of a request, including headers and query strings, but not including the body content | Each supported Region: 20,480 Bytes | No | The maximum length of a request (in bytes), including headers and query strings, but not including the body content. |
| Maximum number of CloudFront distributions that can be associated with a field-level encryption configuration | Each supported Region: 20 | No | The maximum number of CloudFront distributions that can be associated with a field-level encryption configuration. |
| Maximum number of aliases per distribution tenant | Each supported Region: 100 | Yes | The maximum number of aliases that you can have for a distribution tenant. |
| Maximum number of associations per web ACL | Each supported Region: 100 | Yes | The maximum number of resources that can be associated with a single web ACL. |
| Maximum number of certificates per trust store | Each supported Region: 25 | Yes | The maximum number of certificates that you can have for a trust store. |
| Maximum number of characters in a whitelisted query string | Each supported Region: 128 | No | The maximum number of characters in a whitelisted query string. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of characters total for all whitelisted query strings in the same parameter | Each supported Region: 512 | No | The maximum number of characters total for all whitelisted query strings in the same parameter. |
| Maximum number of connection functions per AWS account | Each supported Region: 0 | Yes | The maximum number of connection functions that you can have for an AWS account. |
| Maximum number of connection groups per AWS account | Each supported Region: 100 | Yes | The maximum number of connection groups that you can have for an AWS account. |
| Maximum number of connection groups per Anycast static IP list | Each supported Region: 5 | Yes | The maximum number of connection groups that can be associated with the same Anycast static IP list. |
| Maximum number of distribution tenants per AWS account | Each supported Region: 10,000 | Yes | The maximum number of distribution tenants that you can have for an AWS account. |
| Maximum number of distributions associated with the same trust store | Each supported Region: 25 | Yes | The maximum number of distributions that can be associated with a single trust store. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of field-level encryption configurations that can be associated with one AWS account | Each supported Region: 10 | No | The maximum number of field-level encryption configurations that can be associated with one AWS account. |
| Maximum number of field-level encryption profiles that can be associated with one AWS account | Each supported Region: 10 | No | The maximum number of field-level encryption profiles that can be associated with one AWS account. |
| Maximum number of fields in a request body when field-level encryption is configured | Each supported Region: 10 | No | The maximum number of fields in a request body when field-level encryption is configured. |
| Maximum number of fields to encrypt that can be specified in one profile | Each supported Region: 10 | No | The maximum number of fields to encrypt that can be specified in one profile. |
| Maximum number of multi-tenant distributions per AWS account | Each supported Region: 20 | Yes | The maximum number of multi-tenant distributions that you can have for an AWS account. |
| Maximum number of parameters in a field in a multi-tenant distribution | Each supported Region: 2 | Yes | The maximum number of parameters that you can have in a field in a multi-tenant distribution. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of parameters per distribution tenant | Each supported Region: 5 | Yes | The maximum number of parameters that you can have for a distribution tenant. |
| Maximum number of parameters per multi-tenant distribution | Each supported Region: 5 | Yes | The maximum number of parameters that you can have for a multi-tenant distribution. |
| Maximum number of public keys that can be added to one AWS account | Each supported Region: 10 | No | The maximum number of public keys that can be added to one AWS account. |
| Maximum number of query argument profile mappings that can be included in a field-level encryption configuration | Each supported Region: 5 | No | The maximum number of query argument profile mappings that can be included in a field-level encryption configuration. |
| Maximum number of trust stores per AWS account | Each supported Region: 20 | Yes | The maximum number of trust stores that you can have for an AWS account. |
| Origin access identities per account | Each supported Region: 100 | Yes | The maximum number of origin access identities per account. |
| Origin groups per distribution | Each supported Region: 10 | Yes | The maximum number of origin groups per distribution. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Origin request policies per AWS account | Each supported Region: 20 | Yes | The maximum number of origin request policies per AWS account. |
| Origin response timeout (idle timeout) | Each supported Region: 10 | No | The maximum origin response timeout (idle timeout) in minutes. If CloudFront hasn't detected any bytes sent from the origin to the client within the past 10 minutes, the connection is assumed to be idle and is closed. |
| Origins per distribution | Each supported Region: 100 | Yes | The maximum number of origins per distribution. |
| Public keys in a single key group | Each supported Region: 5 | Yes | The maximum number of public keys in a single key group. |
| Query strings per cache policy | Each supported Region: 10 | Yes | The maximum number of query strings per cache policy. |
| Query strings per origin request policy | Each supported Region: 10 | Yes | The maximum number of query strings per origin request policy. |
| Range of file sizes that CloudFront compresses | Each supported Region: 10,000,000 Bytes | No | The range of file sizes (in bytes) that CloudFront compresses (1,000 to 10,000,000). |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Removing headers from a response headers policy | Each supported Region: 10 | Yes | The maximum number of headers that you can remove in a response headers policy. |
| Request body size for origin requests exposed to a Lambda@Edge function. | Each supported Region: 1 Megabytes | No | The maximum request body size (in MB) for origin requests exposed to a Lambda@Edge function. |
| Request body size for origin requests when returning from a Lambda function (base64 encoding) | Each supported Region: 1.33 Megabytes | No | The maximum request body size (in KB) for origin requests when returning from a Lambda function. (base64 encoding) |
| Request body size for origin requests when returning from a Lambda function (text encoding) | Each supported Region: 1 Megabytes | No | The maximum request body size (in KB) for origin requests when returning from a Lambda function. (text encoding) |
| Request body size for viewer requests exposed to a Lambda@Edge function. | Each supported Region: 40 Kilobytes | No | The maximum request body size (in KB) for viewer requests exposed to a Lambda@Edge function. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request body size for viewer requests when returning from a Lambda function (base64 encoding) | Each supported Region: 53.2 Kilobytes | No | The maximum request body size (in KB) for viewer requests when returning from a Lambda function. (base64 encoding) |
| Request body size for viewer requests when returning from a Lambda function (text encoding) | Each supported Region: 40 Kilobytes | No | The maximum request body size (in KB) for viewer requests when returning from a Lambda function. (text encoding) |
| Requests per second per distribution | Each supported Region: 250,000 | Yes | The maximum number of requests per second per distribution. |
| Response headers policies | Each supported Region: 20 | Yes | The maximum number of response headers policies per AWS account. |
| Response timeout per origin | Each supported Region: 120 Seconds | Yes | The response timeout per origin (1-120 seconds). |
| SSL certificates per AWS account when serving HTTPS requests using dedicated IP addresses | Each supported Region: 2 | Yes | The maximum number of SSL certificates per AWS account when serving HTTPS requests using dedicated IP addresses (no quota when serving HTTPS requests using SNI). |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SSL certificates that can be associated with a CloudFront web distribution | Each supported Region: 1 | No | The maximum number of SSL certificates that can be associated with a CloudFront web distribution. |
| Size of a response that is generated by a Lambda function, including headers and body (Origin request and response event) | Each supported Region: 1 Megabytes | No | The maximum size (in MB) of a response that is generated by a Lambda function, including headers and body. (Origin request and response event) |
| Size of a response that is generated by a Lambda function, including headers and body (Viewer request and response event) | Each supported Region: 40 Kilobytes | No | The maximum size (in KB) of a response that is generated by a Lambda function, including headers and body. (Viewer request and response event) |
| Staging distributions per AWS account | Each supported Region: 20 | Yes | The maximum number of staging distributions that you can have for an AWS account. |
| Tags that can be added to a distribution | Each supported Region: 50 | No | The maximum number of tags that can be added to a distribution. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Total length of the URI including query string in a Lambda@Edge function | Each supported Region: 8,192 | No | The maximum total length in characters of the URI including the query string in a Lambda@Edge function. |
| Total number of bytes in whitelisted cookie names (doesn't apply if you configure CloudFront to forward all cookies to the origin) | Each supported Region: 512 Bytes | No | The total number of bytes in whitelisted cookie names (doesn't apply if you configure CloudFront to forward all cookies to the origin). 512 minus the number of whitelisted cookies. |
| Triggers per distribution | Each supported Region: 100 | Yes | The maximum number of triggers per distribution. |
| VPC origins per AWS account | Each supported Region: 25 | Yes | The maximum number of VPC origins per AWS account. |
| Web distributions per AWS account | Each supported Region: 500 | Yes | The maximum number of web distributions per AWS account. |
| Whitelisted cookies per cache behavior | Each supported Region: 10 | Yes | The maximum number of whitelisted cookies per cache behavior. |
| Whitelisted headers per cache behavior | Each supported Region: 10 | Yes | The maximum number of whitelisted headers per cache behavior. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Whitelisted query strings per cache behavior | Each supported Region: 10 | Yes | The maximum number of whitelisted query strings per cache behavior. |

For more information, see Quotas in the *Amazon CloudFront Developer Guide*.

# AWS CloudHSM endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | cloudhsmv2.us-east-2.amazonaws.com | HTTPS |
| | | cloudhsmv2.us-east-2.api.aws | HTTPS |
| | | cloudhsmv2-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | cloudhsmv2.us-east-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.us-east-1.api.aws | HTTPS |
| | | cloudhsmv2-fips.us-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | cloudhsmv2.us-west-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.us-west-1.api.aws | HTTPS |
| | | cloudhsmv2-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | cloudhsmv2.us-west-2.amazonaws.com | HTTPS |
| | | cloudhsmv2.us-west-2.api.aws | HTTPS |
| | | cloudhsmv2-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | cloudhsmv2.af-south-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | cloudhsmv2.ap-east-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | cloudhsmv2.ap-south-2.amazonaws.com | HTTPS |
| | | cloudhsmv2.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | cloudhsmv2.ap-southeast-3.amazonaws.com | HTTPS |
| | | cloudhsmv2.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | cloudhsmv2.ap-southeast-5.amazonaws.com | HTTPS |
| | | cloudhsmv2.ap-southeast-5.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | cloudhsmv2.ap-south-1.amazonaws.com<br>cloudhsmv2.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | cloudhsmv2.ap-southeast-6.amazonaws.com<br>cloudhsmv2.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | cloudhsmv2.ap-northeast-3.amazonaws.com<br>cloudhsmv2.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | cloudhsmv2.ap-northeast-2.amazonaws.com<br>cloudhsmv2.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudhsmv2.ap-southeast-1.amazonaws.com<br>cloudhsmv2.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudhsmv2.ap-southeast-2.amazonaws.com<br>cloudhsmv2.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | cloudhsmv2.ap-southeast-7.amazonaws.com<br>cloudhsmv2.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudhsmv2.ap-northeast-1.amazonaws.com<br>cloudhsmv2.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | cloudhsmv2.ca-central-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.ca-central-1.api.aws | HTTPS |
| | | cloudhsmv2-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | cloudhsmv2.ca-west-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.ca-west-1.api.aws | HTTPS |
| | | cloudhsmv2-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | cloudhsmv2.eu-central-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | cloudhsmv2.eu-west-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | cloudhsmv2.eu-west-2.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | cloudhsmv2.eu-south-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | cloudhsmv2.eu-west-3.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | cloudhsmv2.eu-south-2.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | cloudhsmv2.eu-north-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | cloudhsmv2.eu-central-2.amazonaws.com | HTTPS |
| | | cloudhsmv2.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | cloudhsmv2.il-central-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | cloudhsmv2.mx-central-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | cloudhsmv2.me-south-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | cloudhsmv2.me-central-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | cloudhsmv2.sa-east-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | cloudhsmv2.us-gov-east-1.amazonaws.com | HTTPS |
| | | cloudhsmv2.us-gov-east-1.api.aws | HTTPS |
| | | cloudhsmv2-fips.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | cloudhsmv2.us-gov-west-1.amazonaws.com | HTTPS |
| | | cloudhsmv2-fips.us-gov-west-1.api.aws | HTTPS |
| | | cloudhsmv2.us-gov-west-1.api.aws | HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Clusters per AWS Region and AWS account | Each supported Region: 4 | Yes | The maximum number of clusters that you can create in this account in the current Region. |
| HSMs per AWS Region and AWS account | Each supported Region: 6 | Yes | The maximum number of HSMs that you can create in this account in the current Region. |
| HSMs per CloudHSM cluster | Each supported Region: 28 | No | The maximum number of HSMs that you can create in a CloudHSM cluster. |
| Keys per CloudHSM cluster | Each supported Region: 3,300 | No | The maximum number of keys that you can create in a CloudHSM cluster. |
| Length of a Username | Each supported Region: 31 | No | The maximum number of characters for a username. |
| Length of a password | Each supported Region: 32 | No | The maximum number of characters for a password. |
| Minimum length of a password | Each supported Region: 7 | No | The minimum number of characters for a password. |
| Number of concurrent clients | Each supported Region: 900 | No | The maximum number of concurrent clients that can exist in a Region. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Users per CloudHSM cluster | Each supported Region: 1,024 | No | The maximum number of users who can be created on a cluster in an account. |

For more information, see Quotas in the *AWS CloudHSM User Guide*.

# AWS Cloud Map endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | servicediscovery.us-east-2.amazonaws.com | HTTPS | |
| | | servicediscovery.us-east-2.api.aws | HTTPS | |
| | | servicediscovery-fips.us-east-2.api.aws | HTTPS | |
| | | servicediscovery-fips.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | servicediscovery.us-east-1.amazonaws.com | HTTPS |
| | | servicediscovery-fips.us-east-1.amazonaws.com | HTTPS |
| | | servicediscovery.us-east-1.api.aws | HTTPS |
| | | servicediscovery-fips.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | servicediscovery.us-west-1.amazonaws.com | HTTPS |
| | | servicediscovery.us-west-1.api.aws | HTTPS |
| | | servicediscovery-fips.us-west-1.api.aws | HTTPS |
| | | servicediscovery-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | servicediscovery.us-west-2.amazonaws.com | HTTPS |
| | | servicediscovery-fips.us-west-2.amazonaws.com | HTTPS |
| | | servicediscovery.us-west-2.api.aws | HTTPS |
| | | servicediscovery-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | servicediscovery.af-south-1.amazonaws.com | HTTPS |
| | | servicediscovery.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | servicediscovery.ap-east-1.amazonaws.com | HTTPS |
| | | servicediscovery.ap-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | servicediscovery.ap-south-2.amazonaws.com<br>servicediscovery.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | servicediscovery.ap-southeast-3.amazonaws.com<br>servicediscovery.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | servicediscovery.ap-southeast-5.amazonaws.com<br>servicediscovery.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | servicediscovery.ap-southeast-4.amazonaws.com<br>servicediscovery.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | servicediscovery.ap-south-1.amazonaws.com<br>servicediscovery.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | servicediscovery.ap-southeast-6.amazonaws.com<br>servicediscovery.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | servicediscovery.ap-northeast-3.amazonaws.com<br>servicediscovery.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | servicediscovery.ap-northeast-2.amazonaws.com<br>servicediscovery.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | servicediscovery.ap-southeast-1.amazonaws.com<br>servicediscovery.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | servicediscovery.ap-southeast-2.amazonaws.com<br>servicediscovery.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | servicediscovery.ap-east-2.amazonaws.com<br>servicediscovery.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | servicediscovery.ap-southeast-7.amazonaws.com<br>servicediscovery.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | servicediscovery.ap-northeast-1.amazonaws.com<br>servicediscovery.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | servicediscovery.ca-central-1.amazonaws.com<br>servicediscovery.ca-central-1.api.aws<br>servicediscovery-fips.ca-central-1.amazonaws.com<br>servicediscovery-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | servicediscovery.ca-west-1.amazonaws.com | HTTPS |
| | | servicediscovery-fips.ca-west-1.amazonaws.com | HTTPS |
| | | servicediscovery.ca-west-1.api.aws | HTTPS |
| | | servicediscovery-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | servicediscovery.eu-central-1.amazonaws.com | HTTPS |
| | | servicediscovery.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | servicediscovery.eu-west-1.amazonaws.com | HTTPS |
| | | servicediscovery.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | servicediscovery.eu-west-2.amazonaws.com | HTTPS |
| | | servicediscovery.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | servicediscovery.eu-south-1.amazonaws.com | HTTPS |
| | | servicediscovery.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | servicediscovery.eu-west-3.amazonaws.com | HTTPS |
| | | servicediscovery.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | servicediscovery.eu-south-2.amazonaws.com | HTTPS |
| | | servicediscovery.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | servicediscovery.eu-north-1.amazonaws.com | HTTPS |
| | | servicediscovery.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | servicediscovery.eu-central-2.amazonaws.com | HTTPS |
| | | servicediscovery.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | servicediscovery.il-central-1.amazonaws.com | HTTPS |
| | | servicediscovery.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | servicediscovery.mx-central-1.amazonaws.com | HTTPS |
| | | servicediscovery.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | servicediscovery.me-south-1.amazonaws.com | HTTPS |
| | | servicediscovery.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | servicediscovery.me-central-1.amazonaws.com | HTTPS |
| | | servicediscovery.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | servicediscovery.sa-east-1.amazonaws.com | HTTPS |
| | | servicediscovery.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | servicediscovery.us-gov-east-1.amazonaws.com | HTTPS |
| | | servicediscovery-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | servicediscovery.us-gov-east-1.api.aws | HTTPS |
| | | servicediscovery-fips.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | servicediscovery.us-gov-west-1.amazonaws.com<br>servicediscovery-fips.us-gov-west-1.api.aws<br>servicediscovery-fips.us-gov-west-1.amazonaws.com<br>servicediscovery.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom attributes per instance | Each supported Region: 30 | No | The maximum number of custom attributes that you can specify when you register an instance. |
| DiscoverInstances operation per account burst rate | Each supported Region: 2,000 | Yes | The maximum burst rate to call DiscoverInstances operation from a single account. |
| DiscoverInstances operation per account steady rate | Each supported Region: 1,000 | Yes | The maximum steady rate to call DiscoverInstances operation from a single account. |
| DiscoverInstancesRevision operation per account rate | Each supported Region: 3,000 | Yes | The maximum rate to call DiscoverInstancesRevision operation from a single account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Instances per namespace | Each supported Region: 2,000 | Yes | The maximum number of service instances that you can register using the same namespace. |
| Instances per service | Each supported Region: 1,000 | No | The maximum number of instances that you can register in a Region using the same service. |
| Namespaces per Region | Each supported Region: 50 | Yes | The maximum number of namespaces that you can create per Region. |

For more information, see AWS Cloud Map Quotas and AWS Cloud Map DiscoverInstances API request throttling in the *AWS Cloud Map Developer Guide*.

## Amazon CloudSearch endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | cloudsearch.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | cloudsearch.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | cloudsearch.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cloudsearch.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudsearch.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudsearch.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudsearch.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cloudsearch.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | cloudsearch.eu-west-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | cloudsearch.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Document batch size | Each supported Region: 5 Megabytes | No | The size of document batch uploads. |
| Document size | Each supported Region: 1 Megabytes | No | The size of individual documents. |
| Domains per account | Each supported Region: 100 | No | Number of search domains you can create per AWS account. |
| Index fields | Each supported Region: 200 | Yes | The number of index fields per domain. A dynamic field counts as one index field, but typically matches multiple document fields. Dynamic fields can cause the total number of fields in your index |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | to exceed this limit. If you use dynamic fields, keep the number of index fields below 1,000 to avoid performance issues. |
| Partition count | Each supported Region: 10 | Yes | The minimum number of partitions your search index is distributed across per Amazon CloudSearch domain. |
| Replication count | Each supported Region: 5 | Yes | The minimum number of replicas each partition of your Amazon CloudSearch domain would contain. |
| Search document fields | Each supported Region: 200 | No | Number of fields per document. |

For more information, see Understanding Amazon CloudSearch Quotas in the *Amazon CloudSearch Developer Guide*.

## AWS CloudShell endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cloudshell.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | cloudshell.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | cloudshell.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | cloudshell.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | cloudshell.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cloudshell.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | cloudshell.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | cloudshell.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | cloudshell.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | cloudshell.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cloudshell.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cloudshell.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cloudshell.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudshell.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudshell.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | cloudshell.ap-southeast-7.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudshell.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | cloudshell.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | cloudshell.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cloudshell.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | cloudshell.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cloudshell.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | cloudshell.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | cloudshell.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | cloudshell.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | cloudshell.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | cloudshell.eu-central-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | cloudshell.il-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | cloudshell.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | cloudshell.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | cloudshell.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | cloudshell.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | cloudshell.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Concurrent shells | Each supported Region: 10 | Yes | The maximum number of concurrent shells that you can run at the same |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| | | | time in this account in the current Region. |
| Data retention | Each supported Region: 120 | No | The number of days that the data in the home directory will be retained after a shell was last accessed. |
| Home directory size | Each supported Region: 1 Gigabytes | No | The maximum size of your shells home directory. |
| Monthly usage | Each supported Region: 200 | Yes | The maximum number of hours that you can use AWS CloudShell per month in this account in the current Region. |

## Amazon CodeCatalyst endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | codecatalyst.global.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | codecatalyst.global.api.aws | HTTPS |

## Service quotas

For information on quotas, see Quotas for Amazon CodeCatalyst in the *Amazon CodeCatalyst User Guide*.

# AWS Control Catalog endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service Endpoints

AWS Control Catalog has the following endpoints:

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | controlcatalog.us-east-2.amazonaws.com | HTTPS | |
| | | controlcatalog-fips.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | controlcatalog-fips.us-east-2.amazonaws.com | HTTPS |
| | | controlcatalog.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | controlcatalog.us-east-1.amazonaws.com | HTTPS |
| | | controlcatalog-fips.us-east-1.api.aws | HTTPS |
| | | controlcatalog-fips.us-east-1.amazonaws.com | HTTPS |
| | | controlcatalog.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | controlcatalog.us-west-1.amazonaws.com | HTTPS |
| | | controlcatalog-fips.us-west-1.api.aws | HTTPS |
| | | controlcatalog-fips.us-west-1.amazonaws.com | HTTPS |
| | | controlcatalog.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | controlcatalog.us-west-2.amazonaws.com | HTTPS |
| | | controlcatalog-fips.us-west-2.api.aws | HTTPS |
| | | controlcatalog-fips.us-west-2.amazonaws.com | HTTPS |
| | | controlcatalog.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | controlcatalog.af-south-1.amazonaws.com | HTTPS |
| | | controlcatalog.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | controlcatalog.ap-east-1.amazonaws.com | HTTPS |
| | | controlcatalog.ap-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | controlcatalog.ap-south-2.amazonaws.com<br>controlcatalog.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | controlcatalog.ap-southeast-3.amazonaws.com<br>controlcatalog.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | controlcatalog.ap-southeast-5.amazonaws.com<br>controlcatalog.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | controlcatalog.ap-southeast-4.amazonaws.com<br>controlcatalog.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | controlcatalog.ap-south-1.amazonaws.com<br>controlcatalog.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | controlcatalog.ap-southeast-6.amazonaws.com<br>controlcatalog.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | controlcatalog.ap-northeast-3.amazonaws.com<br>controlcatalog.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | controlcatalog.ap-northeast-2.amazonaws.com<br>controlcatalog.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | controlcatalog.ap-southeast-1.amazonaws.com<br>controlcatalog.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | controlcatalog.ap-southeast-2.amazonaws.com<br>controlcatalog.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | controlcatalog.ap-east-2.amazonaws.com<br>controlcatalog.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | controlcatalog.ap-southeast-7.amazonaws.com<br>controlcatalog.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | controlcatalog.ap-northeast-1.amazonaws.com<br>controlcatalog.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | controlcatalog.ca-central-1.amazonaws.com<br>controlcatalog-fips.ca-central-1.api.aws<br>controlcatalog-fips.ca-central-1.amazonaws.com<br>controlcatalog.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | controlcatalog.ca-west-1.amazonaws.com | HTTPS |
| | | controlcatalog-fips.ca-west-1.api.aws | HTTPS |
| | | controlcatalog-fips.ca-west-1.amazonaws.com | HTTPS |
| | | controlcatalog.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | controlcatalog.eu-central-1.amazonaws.com | HTTPS |
| | | controlcatalog.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | controlcatalog.eu-west-1.amazonaws.com | HTTPS |
| | | controlcatalog.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | controlcatalog.eu-west-2.amazonaws.com | HTTPS |
| | | controlcatalog.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | controlcatalog.eu-south-1.amazonaws.com | HTTPS |
| | | controlcatalog.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | controlcatalog.eu-west-3.amazonaws.com | HTTPS |
| | | controlcatalog.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | controlcatalog.eu-south-2.amazonaws.com | HTTPS |
| | | controlcatalog.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | controlcatalog.eu-north-1.amazonaws.com | HTTPS |
| | | controlcatalog.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | controlcatalog.eu-central-2.amazonaws.com | HTTPS |
| | | controlcatalog.eu-central-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | controlcatalog.il-central-1.amazonaws.com<br>controlcatalog.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | controlcatalog.mx-central-1.amazonaws.com<br>controlcatalog.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | controlcatalog.me-south-1.amazonaws.com<br>controlcatalog.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | controlcatalog.me-central-1.amazonaws.com<br>controlcatalog.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | controlcatalog.sa-east-1.amazonaws.com<br>controlcatalog.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | controlcatalog.us-gov-east-1.amazonaws.com<br>controlcatalog.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | controlcatalog.us-gov-west-1.amazonaws.com<br>controlcatalog.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

## Service Quotas

This service has no quotas.

# AWS CloudTrail endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Control plane endpoints

The following table contains AWS Region-specific endpoints that AWS CloudTrail supports for control plane operations. For more information, see the AWS CloudTrail API Reference.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | cloudtrail.us-east-2.amazonaws.com | HTTPS |
| | | cloudtrail-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | cloudtrail.us-east-1.amazonaws.com | HTTPS |
| | | cloudtrail-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | cloudtrail.us-west-1.amazonaws.com | HTTPS |
| | | cloudtrail-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | cloudtrail.us-west-2.amazonaws.com | HTTPS |
| | | cloudtrail-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | cloudtrail.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cloudtrail.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | cloudtrail.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | cloudtrail.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | cloudtrail.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | cloudtrail.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cloudtrail.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | cloudtrail.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cloudtrail.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cloudtrail.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudtrail.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudtrail.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | cloudtrail.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | cloudtrail.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudtrail.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | cloudtrail.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | cloudtrail.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cloudtrail.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | cloudtrail.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cloudtrail.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | cloudtrail.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | cloudtrail.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | cloudtrail.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | cloudtrail.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | cloudtrail.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | cloudtrail.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | cloudtrail.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | cloudtrail.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | cloudtrail.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | cloudtrail.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | cloudtrail.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | cloudtrail.us-gov-west-1.amazonaws.com | HTTPS | |

## Data plane endpoints

The following table contains AWS Region-specific endpoints that AWS CloudTrail supports for data plane operations. For more information, see the AWS CloudTrail Data API Reference.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cloudtrail-data.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | cloudtrail-data.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | cloudtrail-data.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | cloudtrail-data.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | cloudtrail-data.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cloudtrail-data.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | cloudtrail-data.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cloudtrail-data.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cloudtrail-data.ap-northeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | cloudtrail-data.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cloudtrail-data.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cloudtrail-data.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cloudtrail-data.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | cloudtrail-data.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cloudtrail-data.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | cloudtrail-data.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cloudtrail-data.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | cloudtrail-data.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | cloudtrail-data.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | cloudtrail-data.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | cloudtrail-data.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | cloudtrail-data.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | cloudtrail-data.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Channels | Each supported Region: 25 | No | This quota applies to channels used for CloudTrail Lake integrations with event sources outside of AWS, and does not apply to service-linked channels. |
| CloudTrail file size sent to Amazon S3 | Each supported Region: 50 Megabytes | No | For both management and data events, CloudTrail sends events to S3 in maximum 50 MB |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | (compressed) ZIP files. If enabled on the trail, log delivery notifications are sent by Amazon SNS after CloudTrail sends ZIP files to S3. |
| Concurrent dashboard refreshes | Each supported Region: 1 | No | The maximum number of ongoing refreshes per dashboard. |
| Concurrent queries | Each supported Region: 10 | No | The maximum number of queued or running queries that you can run simultaneously in CloudTrail Lake. |
| Conditions across all advanced event selectors | Each supported Region: 500 | No | If a trail uses advanced event selectors, a maximum of 500 total values for all conditions in all advanced event selectors is allowed. Unless a trail logs data events on all resources, such as all S3 buckets, a trail is limited to 250 data resources. Data resources can be distributed across event selectors, but the total cannot exceed 250. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Custom Dashboards per region | Each supported Region: 100 | No | The maximum number of custom dashboards per region. |
| Data resources across all event selectors in a trail | Each supported Region: 250 | No | If you choose to limit data events by using event selectors or advanced event selectors, the total number of data resources cannot exceed 250 across all event selectors in a trail. |
| Event data stores | Each supported Region: 10 | No | The maximum number of event data stores per region. This includes single-Region event data stores for the Region, multi-Region event data stores across all AWS Regions, and organization event data stores. In shadow regions, to get the latest resource count metric, call the ListEventDataStores API. |
| Event selectors | Each supported Region: 5 | No | The maximum number of event selectors per trail. |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Event size | Each supported Region: 256 Kilobytes | No | The maximum event size (in KB). All event versions: events over 256 KB cannot be sent to CloudWatch Logs. Event version 1.05 and newer: maximum event size of 256 KB. |
| Events per PutAuditEvents request | Each supported Region: 100 | No | You can add up to 100 activity events (or up to 1 MB) per PutAuditEvents request. |
| Summarized query results size | Each supported Region: 3 Megabytes | No | The maximum size (in MB) of query results that can be summarized in a month |
| Trails per region | Each supported Region: 5 | No | The maximum number of trails per region. In shadow regions, to get the latest resource count metric, call the ListTrails API. |
| Transactions per minute (TPM) for the GenerateQuery API | Each supported Region: 4 per minute | No | The maximum number of operation requests you can make per minute without being throttled. |
| Transactions per second (TPS) for the AddTags API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the CancelQuery API | Each supported Region: 3 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the CreateChannel API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the CreateDashboard API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the CreateEventDataStore API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the CreateTrail API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DeleteChannel API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DeleteDashboard API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the DeleteEventDataStore API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DeleteResourcePolicy API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DeleteTrail API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DeregisterOrganizationDelegatedAdmin API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DescribeQuery API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DescribeTrails API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the DisableFederation API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the EnableFederation API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetChannel API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetDashboard API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetEventDataStore API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetEventSelectors API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetImport API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetInsightSelectors API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the GetQueryResults API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetResourcePolicy API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetTrail API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the GetTrailStatus API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListChannels API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListDashboards API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListEventDataStores API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the ListImportFailures API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListImports API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListInsightsMetricData API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListPublicKeys API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListQueries API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListTags API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the ListTrails API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the LookupEvents API | Each supported Region: 2 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the PutAuditEvents API | Each supported Region: 100 | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the PutEventSelectors API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the PutInsightSelectors API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the PutResourcePolicy API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the RegisterOrganizationDelegatedAdmin API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the RemoveTags API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the RestoreEventDataStore API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the SearchSampleQueries API | Each supported Region: 10 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StartDashboardRefresh API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StartEventDataStoreIngestion API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StartImport API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StartLogging API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StartQuery API | Each supported Region: 3 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the StopEventDataStoreIngestion API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StopImport API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the StopLogging API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the UpdateChannel API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the UpdateDashboard API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the UpdateEventDataStore API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |
| Transactions per second (TPS) for the UpdateTrail API | Each supported Region: 1 per second | No | The maximum number of operation requests you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Widgets per dashboard | Each supported Region: 10 | No | The maximum number of widgets per dashboard. |

For more information, see Quotas in AWS CloudTrail.

# Amazon CloudWatch endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | monitoring.us-east-2.amazonaws.com | HTTP and HTTPS | |
| | | monitoring-fips.us-east-2.amazonaws.com | HTTPS | |
| | | monitoring-fips.us-east-2.api.aws | HTTPS | |
| | | monitoring.us-east-2.api.aws | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | monitoring.us-east-1.amazonaws.com<br>monitoring-fips.us-east-1.amazonaws.com<br>monitoring-fips.us-east-1.api.aws<br>monitoring.us-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (N. California) | us-west-1 | monitoring.us-west-1.amazonaws.com<br>monitoring-fips.us-west-1.amazonaws.com<br>monitoring-fips.us-west-1.api.aws<br>monitoring.us-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | monitoring.us-west-2.amazonaws.com<br>monitoring-fips.us-west-2.amazonaws.com<br>monitoring-fips.us-west-2.api.aws<br>monitoring.us-west-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | monitoring.af-south-1.amazonaws.com<br>monitoring.af-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | monitoring.ap-east-1.amazonaws.com<br>monitoring.ap-east-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | monitoring.ap-south-2.amazonaws.com<br>monitoring.ap-south-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | monitoring.ap-southeast-3.amazonaws.com<br>monitoring.ap-southeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | monitoring.ap-southeast-5.amazonaws.com<br>monitoring.ap-southeast-5.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | monitoring.ap-southeast-4.amazonaws.com<br>monitoring.ap-southeast-4.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | monitoring.ap-south-1.amazonaws.com<br>monitoring.ap-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | monitoring.ap-southeast-6.amazonaws.com<br>monitoring.ap-southeast-6.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | monitoring.ap-northeast-3.amazonaws.com<br>monitoring.ap-northeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | monitoring.ap-northeast-2.amazonaws.com<br>monitoring.ap-northeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | monitoring.ap-southeast-1.amazonaws.com<br>monitoring.ap-southeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | monitoring.ap-southeast-2.amazonaws.com<br>monitoring.ap-southeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | monitoring.ap-east-2.amazonaws.com<br>monitoring.ap-east-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | monitoring.ap-southeast-7.amazonaws.com<br>monitoring.ap-southeast-7.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | monitoring.ap-northeast-1.amazonaws.com<br>monitoring.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | monitoring.ca-central-1.amazonaws.com<br>monitoring.ca-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | monitoring.ca-west-1.amazonaws.com<br>monitoring.ca-west-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | monitoring.eu-central-1.amazonaws.com<br>monitoring.eu-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | monitoring.eu-west-1.amazonaws.com<br>monitoring.eu-west-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (London) | eu-west-2 | monitoring.eu-west-2.amazonaws.com<br>monitoring.eu-west-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | monitoring.eu-south-1.amazonaws.com<br>monitoring.eu-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | monitoring.eu-west-3.amazonaws.com<br>monitoring.eu-west-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | monitoring.eu-south-2.amazonaws.com<br>monitoring.eu-south-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | monitoring.eu-north-1.amazonaws.com<br>monitoring.eu-north-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | monitoring.eu-central-2.amazonaws.com<br>monitoring.eu-central-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Israel (Tel Aviv) | il-central-1 | monitoring.il-central-1.amazonaws.com<br>monitoring.il-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Mexico (Central) | mx-central-1 | monitoring.mx-central-1.amazonaws.com<br>monitoring.mx-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | monitoring.me-south-1.amazonaws.com<br>monitoring.me-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | monitoring.me-central-1.amazonaws.com<br><br>monitoring.me-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | monitoring.sa-east-1.amazonaws.com<br><br>monitoring.sa-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | monitoring.us-gov-east-1.amazonaws.com<br><br>monitoring.us-gov-east-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | monitoring.us-gov-west-1.amazonaws.com<br><br>monitoring.us-gov-west-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Canary limit | us-east-1: 300<br>ap-east-2: 200<br>ap-northeast-1: 300<br>Each of the other supported Regions: 500 | Yes | The maximum number of canaries per account per region. |
| Minimum frequency | Each supported Region: 60,000 Milliseconds | No | The minimum time, in milliseconds, between runs of the same canary. |
| Number of Alarm Mute Rules | Each supported Region: 2,000 | No | The maximum number of Alarm Mute Rules that you can have in this account in the current region. |
| Number of Contributor Insights rules | Each supported Region: 100 | Yes | The maximum number of Contributor Insights rules you can have in this account. |
| Number of Metrics Insights alarms | Each supported Region: 200 | No | The maximum number of Metrics Insights alarms that you can have in this account in the current region. |
| Rate of ANOMALY_DETECTION_BAND usage in GetMetricData | Each supported Region: 1,000 | No | The maximum number of times the ANOMALY_D |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | ETECTION_BAND function can be used in all GetMetricData requests, per second, in this account in the current region. |
| Rate of DB_PERF_INSIGHTS usage in GetMetricData | Each supported Region: 4 | No | The maximum number of times the DB_PERF_INSIGHTS function can be used in all GetMetricData requests, per second, in this account in the current region. |
| Rate of DeleteAlarmMuteRule requests | Each supported Region: 3 per second | No | The maximum number of DeleteAlarmMuteRule requests that you can make, per second, in this account in the current region. |
| Rate of DeleteAlarms requests | Each supported Region: 3 per second | No | The maximum number of DeleteAlarms requests that you can make, per second, in this account in the current region. |
| Rate of DeleteAnomalyDetector requests | Each supported Region: 5 per second | No | The maximum number of DeleteAnomalyDetector requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteDashboards requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteDashboards requests that you can make, per second, in this account in the current region. |
| Rate of DeleteInsightRules requests | Each supported Region: 5 per second | No | The maximum number of DeleteInsightRules requests you can make per second in this account. |
| Rate of DeleteMetricStream requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteMetricStream requests that you can make, per second, in this account in the current region. |
| Rate of DescribeAlarmContributors requests | Each supported Region: 9 per second | Yes | The maximum number of DescribeAlarmContributors requests that you can make, per second, in this account in the current region. |
| Rate of DescribeAlarmHistory requests | Each supported Region: 20 per second | No | The maximum number of DescribeAlarmHistory requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeAlarms requests | Each supported Region: 9 per second | Yes | The maximum number of DescribeAlarms requests that you can make, per second, in this account in the current region. |
| Rate of DescribeAlarmsForMetric requests | Each supported Region: 9 per second | No | The maximum number of DescribeAlarmsForMetric requests that you can make, per second, in this account in the current region. |
| Rate of DescribeAnomalyDetectors requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeAnomalyDetectors requests that you can make, per second, in this account in the current region. |
| Rate of DescribeInsightRules requests | Each supported Region: 20 per second | No | The maximum number of DescribeInsightRules requests you can make per second in this account. |
| Rate of DisableAlarmActions requests | Each supported Region: 3 per second | No | The maximum number of DisableAlarmActions requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DisableInsightRules requests | Each supported Region: 1 per second | No | The maximum number of DisableInsightRules requests you can make per second in this account. |
| Rate of EnableAlarmActions requests | Each supported Region: 3 per second | No | The maximum number of EnableAlarmActions requests that you can make, per second, in this account in the current region. |
| Rate of EnableInsightRules requests | Each supported Region: 1 per second | No | The maximum number of EnableInsightRules requests you can make per second in this account. |
| Rate of GetAlarmMuteRule requests | Each supported Region: 20 per second | Yes | The maximum number of GetAlarmMuteRule requests that you can make, per second, in this account in the current region. |
| Rate of GetDashboard requests | Each supported Region: 10 per second | Yes | The maximum number of GetDashboard requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetInsightRuleReport requests | Each supported Region: 20 per second | Yes | The maximum number of GetInsightRuleReport requests you can make, per second in this account. |
| Rate of GetMetricData datapoints for metrics older than three hours | Each supported Region: 396,000 | No | The maximum number of GetMetricData datapoints that you can fetch, per second, for a request with a StartTime of more than three hours in this account in the current region. |
| Rate of GetMetricData datapoints for the last three hours of metrics | Each supported Region: 180,000 | No | The maximum number of GetMetricData datapoints that you can fetch, per second, for a request with a StartTime of less than or equal to three hours in this account in the current region. |
| Rate of GetMetricData datapoints using Metrics Insights | Each supported Region: 4,300,000 | No | The maximum number of GetMetricData datapoints that you can fetch using Metrics Insights, per second, for a request with a StartTime of less than or equal to three hours in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetMetricData requests | Each supported Region: 500 per second | Yes | The maximum number of GetMetricData requests that you can make, per second, in this account in the current region. |
| Rate of GetMetricStatistics requests | Each supported Region: 400 per second | Yes | The maximum number of GetMetricStatistics requests that you can make, per second, in this account in the current region. |
| Rate of GetMetricStream requests | Each supported Region: 10 per second | Yes | The maximum number of GetMetricStream requests that you can make, per second, in this account in the current region. |
| Rate of GetMetricWidgetImage requests | Each supported Region: 20 per second | Yes | The maximum number of GetMetricWidgetImage requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of INSIGHT_RULE_METRIC usage in GetMetricData | Each supported Region: 20 | No | The maximum number of times the INSIGHT_RULE_METRIC function can be used in all GetMetricData requests, per second, in this account in the current region. |
| Rate of LAMBDA usage in GetMetricData | Each supported Region: 5 | No | The maximum number of times the LAMBDA function can be used in all GetMetricData requests, per second, in this account in the current region. |
| Rate of ListAlarmMuteRules requests | Each supported Region: 50 per second | Yes | The maximum number of ListAlarmMuteRules requests that you can make, per second, in this account in the current region. |
| Rate of ListDashboards requests | Each supported Region: 10 per second | Yes | The maximum number of ListDashboards requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListMetricStreams requests | Each supported Region: 10 per second | Yes | The maximum number of ListMetricStreams requests that you can make, per second, in this account in the current region. |
| Rate of ListMetrics requests | Each supported Region: 25 per second | Yes | The maximum number of ListMetrics requests that you can make, per second, in this account in the current region. |
| Rate of ListTagsForResource requests | Each supported Region: 10 per second | No | The maximum number of ListTagsForResource requests that you can make, per second, in this account in the current region. |
| Rate of Metrics Insights usage in GetMetricData | Each supported Region: 10 | No | The maximum number of times Metrics Insights can be used in all GetMetricData requests, per second, in this account in the current region. |
| Rate of PutAlarmMuteRule requests | Each supported Region: 3 per second | Yes | The maximum number of PutAlarmMuteRule requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of PutAnomalyDetector requests | Each supported Region: 10 per second | Yes | The maximum number of PutAnomalyDetector requests that you can make, per second, in this account in the current region. |
| Rate of PutCompositeAlarm requests | Each supported Region: 3 per second | Yes | The maximum number of PutCompositeAlarm requests that you can make, per second, in this account in the current region. |
| Rate of PutDashboard requests | Each supported Region: 10 per second | Yes | The maximum number of PutDashboard requests that you can make, per second, in this account in the current region. |
| Rate of PutInsightRule requests | Each supported Region: 5 per second | No | The maximum number of PutInsightRule requests you can make, per second in this account. |
| Rate of PutLogAlarm requests | Each supported Region: 3 per second | No | The maximum number of PutLogAlarm requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of PutMetricAlarm requests | Each supported Region: 3 per second | Yes | The maximum number of PutMetricAlarm requests that you can make, per second, in this account in the current region. |
| Rate of PutMetricData requests | Each supported Region: 500 per second | Yes | The maximum number of PutMetricData requests that you can make, per second, in this account in the current region. |
| Rate of PutMetricStream requests | Each supported Region: 10 per second | Yes | The maximum number of PutMetricStream requests that you can make, per second, in this account in the current region. |
| Rate of SEARCH usage in GetMetricData | Each supported Region: 50 | No | The maximum number of times the SEARCH function can be used in all GetMetricData requests, per second, in this account in the current region. |
| Rate of SERVICE_QUOTA usage in GetMetricData | Each supported Region: 1,000 | No | The maximum number of times the SERVICE_QUOTA function can be used in all GetMetricData requests, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of SetAlarmState requests | Each supported Region: 3 per second | No | The maximum number of SetAlarmState requests that you can make, per second, in this account in the current region. |
| Rate of StartMetricStreams requests | Each supported Region: 10 per second | Yes | The maximum number of StartMetricStreams requests that you can make, per second, in this account in the current region. |
| Rate of StopMetricStreams requests | Each supported Region: 10 per second | Yes | The maximum number of StopMetricStreams requests that you can make, per second, in this account in the current region. |
| Rate of TagResource requests | Each supported Region: 20 per second | No | The maximum number of TagResource requests that you can make, per second, in this account in the current region. |
| Rate of UntagResource requests | Each supported Region: 20 per second | No | The maximum number of UntagResource requests that you can make, per second, in this account in the current region. |

For more information, see CloudWatch Quotas in the *Amazon CloudWatch User Guide*.

# Amazon CloudWatch Application Insights endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | applicationinsights.us-east-2.amazonaws.com | HTTPS |
| | | applicationinsights-fips.us-east-2.api.aws | HTTPS |
| | | applicationinsights-fips.us-east-2.amazonaws.com | HTTPS |
| | | applicationinsights.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | applicationinsights.us-east-1.amazonaws.com | HTTPS |
| | | applicationinsights-fips.us-east-1.api.aws | HTTPS |
| | | applicationinsights-fips.us-east-1.amazonaws.com | HTTPS |
| | | applicationinsights.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | applicationinsights.us-west-1.amazonaws.com | HTTPS |
| | | applicationinsights-fips.us-west-1.api.aws | HTTPS |
| | | applicationinsights-fips.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | applicationinsights.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | applicationinsights.us-west-2.amazonaws.com<br>applicationinsights-fips.us-west-2.api.aws<br>applicationinsights-fips.us-west-2.amazonaws.com<br>applicationinsights.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | applicationinsights.af-south-1.amazonaws.com<br>applicationinsights.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | applicationinsights.ap-east-1.amazonaws.com<br>applicationinsights.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | applicationinsights.ap-south-2.amazonaws.com<br>applicationinsights.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | applicationinsights.ap-southeast-3.amazonaws.com<br>applicationinsights.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | applicationinsights.ap-southeast-4.amazonaws.com<br>applicationinsights.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | applicationinsights.ap-south-1.amazonaws.com<br><br>applicationinsights.ap-south-1.api.aws | HTTPS<br><br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | applicationinsights.ap-northeast-3.amazonaws.com<br><br>applicationinsights.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | applicationinsights.ap-northeast-2.amazonaws.com<br><br>applicationinsights.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | applicationinsights.ap-southeast-1.amazonaws.com<br><br>applicationinsights.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | applicationinsights.ap-southeast-2.amazonaws.com<br><br>applicationinsights.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | applicationinsights.ap-northeast-1.amazonaws.com<br><br>applicationinsights.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | applicationinsights.ca-central-1.amazonaws.com<br>applicationinsights-fips.ca-central-1.api.aws<br>applicationinsights-fips.ca-central-1.amazonaws.com<br>applicationinsights.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | applicationinsights.ca-west-1.amazonaws.com<br>applicationinsights-fips.ca-west-1.api.aws<br>applicationinsights-fips.ca-west-1.amazonaws.com<br>applicationinsights.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | applicationinsights.eu-central-1.amazonaws.com<br>applicationinsights.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | applicationinsights.eu-west-1.amazonaws.com<br>applicationinsights.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | applicationinsights.eu-west-2.amazonaws.com<br>applicationinsights.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | applicationinsights.eu-south-1.amazonaws.com<br>applicationinsights.eu-south-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | applicationinsights.eu-west-3.amazonaws.com<br>applicationinsights.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | applicationinsights.eu-south-2.amazonaws.com<br>applicationinsights.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | applicationinsights.eu-north-1.amazonaws.com<br>applicationinsights.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | applicationinsights.eu-central-2.amazonaws.com<br>applicationinsights.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | applicationinsights.il-central-1.amazonaws.com<br>applicationinsights.il-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | applicationinsights.me-south-1.amazonaws.com<br>applicationinsights.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | applicationinsights.me-central-1.amazonaws.com<br>applicationinsights.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | applicationinsights.sa-east-1.amazonaws.com<br>applicationinsights.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | applicationinsights.us-gov-east-1.amazonaws.com<br>applicationinsights-fips.us-gov-east-1.api.aws<br>applicationinsights-fips.us-gov-east-1.amazonaws.com<br>applicationinsights.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | applicationinsights.us-gov-west-1.amazonaws.com<br>applicationinsights-fips.us-gov-west-1.api.aws<br>applicationinsights-fips.us-gov-west-1.amazonaws.com<br>applicationinsights.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Resource | Default quota |
|---|---|
| API requests | All API actions are throttled to 5 TPS |
| Resource Group applications | 100 per account |
| Account applications | 1 per account |
| Log Streams | 5 per resource |
| Observations per problem | 20 per dashboard<br>40 per DescribeProblemObservations action |

| Resource | Default quota |
| --- | --- |
| Metrics | 60 per resource |

# Amazon CloudWatch Application Signals endpoints and quotas

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | application-signals.us-east-2.api.aws<br>application-signals-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | application-signals.us-east-1.api.aws<br>application-signals-fips.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | application-signals.us-west-1.api.aws<br>application-signals-fips.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | application-signals.us-west-2.api.aws<br>application-signals-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | application-signals.af-south-1.api.aws | | |
| Asia Pacific (Hong Kong) | ap-east-1 | application-signals.ap-east-1.api.aws | | |
| Asia Pacific | ap-south-2 | application-signals.ap-south-2.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | application-signals.ap-southeast-3.api.aws | | |
| Asia Pacific (Malaysia) | ap-southeast-5 | application-signals.ap-southeast-5.api.aws | | |
| Asia Pacific (Melbourne) | ap-southeast-4 | application-signals.ap-southeast-4.api.aws | | |
| Asia Pacific (Mumbai) | ap-south-1 | application-signals.ap-south-1.api.aws | | |
| Asia Pacific (New Zealand) | ap-southeast-6 | application-signals.ap-southeast-6.api.aws | | |
| Asia Pacific (Osaka) | ap-northeast-3 | application-signals.ap-northeast-3.api.aws | | |
| Asia Pacific (Seoul) | ap-northeast-2 | application-signals.ap-northeast-2.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | application-signals.ap-southeast-1.api.aws | | |
| Asia Pacific (Sydney) | ap-southeast-2 | application-signals.ap-southeast-2.api.aws | | |
| Asia Pacific (Taipei) | ap-east-2 | application-signals.ap-east-2.api.aws | | |
| Asia Pacific (Thailand) | ap-southeast-7 | application-signals.ap-southeast-7.api.aws | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | application-signals.ap-northeast-1.api.aws | | |
| Canada (Central) | ca-central-1 | application-signals.ca-central-1.api.aws<br>application-signals-fips.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | application-signals.ca-west-1.api.aws<br>application-signals-fips.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | application-signals.eu-central-1.api.aws | | |
| Europe (Ireland) | eu-west-1 | application-signals.eu-west-1.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | application-signals.eu-west-2.api.aws | | |
| Europe (Milan) | eu-south-1 | application-signals.eu-south-1.api.aws | | |
| Europe (Paris) | eu-west-3 | application-signals.eu-west-3.api.aws | | |
| Europe (Spain) | eu-south-2 | application-signals.eu-south-2.api.aws | | |
| Europe (Stockholm) | eu-north-1 | application-signals.eu-north-1.api.aws | | |
| Europe (Zurich) | eu-central-2 | application-signals.eu-central-2.api.aws | | |
| Israel (Tel Aviv) | il-central-1 | application-signals.il-central-1.api.aws | | |
| Mexico (Central) | mx-central-1 | application-signals.mx-central-1.api.aws | | |
| Middle East (Bahrain) | me-south-1 | application-signals.me-south-1.api.aws | | |
| Middle East (UAE) | me-central-1 | application-signals.me-central-1.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | application-signals.sa-east-1.api.aws | | |
| AWS GovCloud (US-East) | us-gov-east-1 | application-signals.us-gov-east-1.api.aws<br>application-signals-fips.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | application-signals.us-gov-west-1.api.aws<br>application-signals-fips.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of SLOs per Region | Each supported Region: 250 | Yes | The maximum number of SLOs that you can have in this account in the current region. |
| Number of SLOs per Service | Each supported Region: 100 | Yes | The maximum number of SLOs that you can have per service in the current region. |
| Rate of BatchGetServiceLevelObjectiveBudgetReport requests | Each supported Region: 10 per second | Yes | The maximum number of BatchGetServiceLevelObjectiveBudgetReport requests you can |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | make per second in this Region. |
| Rate of BatchUpdateExclusionWindows requests | Each supported Region: 5 per second | Yes | The maximum number of BatchUpdateExclusionWindows requests you can make per second in this Region. |
| Rate of CreateServiceLevelObjective requests | Each supported Region: 5 per second | Yes | The maximum number of CreateServiceLevelObjective requests you can make per second in this Region. |
| Rate of DeleteGroupingConfiguration requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteGroupingConfiguration requests you can make per second in this Region. |
| Rate of DeleteServiceLevelObjective requests | Each supported Region: 5 per second | Yes | The maximum number of DeleteServiceLevelObjective requests you can make per second in this Region. |
| Rate of GetService requests | Each supported Region: 10 per second | Yes | The maximum number of GetService requests you can make per second in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetServiceLevelObjective requests | Each supported Region: 5 per second | Yes | The maximum number of GetServiceLevelObjective requests you can make per second in this Region. |
| Rate of ListAuditFindings requests | Each supported Region: 1 per second | Yes | The maximum number of ListAuditFindings requests you can make per second in this Region. |
| Rate of ListEntityEvents requests | Each supported Region: 10 per second | Yes | The maximum number of ListEntityEvents requests you can make per second in this Region. |
| Rate of ListGroupingAttributeDefinitions requests | Each supported Region: 10 per second | Yes | The maximum number of ListGroupingAttributeDefinitions requests you can make per second in this Region. |
| Rate of ListServiceDependencies requests | Each supported Region: 10 per second | Yes | The maximum number of ListServiceDependencies requests you can make per second in this Region. |
| Rate of ListServiceDependents requests | Each supported Region: 10 per second | Yes | The maximum number of ListServiceDependents requests you can make per second in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListServiceLevelObjectiveExclusionWindows requests | Each supported Region: 5 per second | Yes | The maximum number of ListServiceLevelObjectiveExclusionWindows requests you can make per second in this Region. |
| Rate of ListServiceLevelObjectives requests | Each supported Region: 5 per second | Yes | The maximum number of ListServiceLevelObjectives requests you can make per second in this Region. |
| Rate of ListServiceOperations requests | Each supported Region: 10 per second | Yes | The maximum number of ListServiceOperations requests you can make per second in this Region. |
| Rate of ListServiceStates requests | Each supported Region: 5 per second | Yes | The maximum number of ListServiceStates requests you can make per second in this Region. |
| Rate of ListServices requests | Each supported Region: 10 per second | Yes | The maximum number of ListServices requests you can make per second in this Region. |
| Rate of ListTagsForResource requests | Each supported Region: 5 per second | Yes | The maximum number of ListTagsForResource requests you can make per second in this Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of PutGroupingConfiguration requests | Each supported Region: 10 per second | Yes | The maximum number of PutGroupingConfiguration requests you can make per second in this Region. |
| Rate of StartDiscovery requests | Each supported Region: 5 per second | Yes | The maximum number of StartDiscovery requests you can make per second in this Region. |
| Rate of TagResource requests | Each supported Region: 5 per second | Yes | The maximum number of TagResource requests you can make per second in this Region. |
| Rate of UntagResource requests | Each supported Region: 5 per second | Yes | The maximum number of UntagResource requests you can make per second in this Region. |
| Rate of UpdateServiceLevelObjective requests | Each supported Region: 5 per second | Yes | The maximum number of UpdateServiceLevelObjective requests you can make per second in this Region. |

For more information, see CloudWatch Quotas in the *Amazon CloudWatch User Guide*.

# Amazon CloudWatch Events endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | events.us-east-2.amazonaws.com | HTTPS |
| | | events-fips.us-east-2.api.aws | HTTPS |
| | | events-fips.us-east-2.amazonaws.com | HTTPS |
| | | events.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | events.us-east-1.amazonaws.com | HTTPS |
| | | events-fips.us-east-1.api.aws | HTTPS |
| | | events-fips.us-east-1.amazonaws.com | HTTPS |
| | | events.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | events.us-west-1.amazonaws.com | HTTPS |
| | | events-fips.us-west-1.api.aws | HTTPS |
| | | events-fips.us-west-1.amazonaws.com | HTTPS |
| | | events.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | events.us-west-2.amazonaws.com | HTTPS |
| | | events-fips.us-west-2.api.aws | HTTPS |
| | | events-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | events.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | events.af-south-1.amazonaws.com | HTTPS |
| | | events.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | events.ap-east-1.amazonaws.com | HTTPS |
| | | events.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | events.ap-south-2.amazonaws.com | HTTPS |
| | | events.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | events.ap-southeast-3.amazonaws.com | HTTPS |
| | | events.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | events.ap-southeast-5.amazonaws.com | HTTPS |
| | | events.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | events.ap-southeast-4.amazonaws.com | HTTPS |
| | | events.ap-southeast-4.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | events.ap-south-1.amazonaws.com | HTTPS |
| | | events.ap-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | events.ap-southeast-6.amazonaws.com<br>events.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | events.ap-northeast-3.amazonaws.com<br>events.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | events.ap-northeast-2.amazonaws.com<br>events.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | events.ap-southeast-1.amazonaws.com<br>events.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | events.ap-southeast-2.amazonaws.com<br>events.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | events.ap-east-2.amazonaws.com<br>events.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | events.ap-southeast-7.amazonaws.com<br>events.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | events.ap-northeast-1.amazonaws.com<br>events.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | events.ca-central-1.amazonaws.com | HTTPS |
| | | events-fips.ca-central-1.api.aws | HTTPS |
| | | events-fips.ca-central-1.amazonaws.com | HTTPS |
| | | events.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | events.ca-west-1.amazonaws.com | HTTPS |
| | | events-fips.ca-west-1.api.aws | HTTPS |
| | | events-fips.ca-west-1.amazonaws.com | HTTPS |
| | | events.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | events.eu-central-1.amazonaws.com | HTTPS |
| | | events.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | events.eu-west-1.amazonaws.com | HTTPS |
| | | events.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | events.eu-west-2.amazonaws.com | HTTPS |
| | | events.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | events.eu-south-1.amazonaws.com | HTTPS |
| | | events.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | events.eu-west-3.amazonaws.com | HTTPS |
| | | events.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | events.eu-south-2.amazonaws.com | HTTPS |
| | | events.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | events.eu-north-1.amazonaws.com<br>events.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | events.eu-central-2.amazonaws.com<br>events.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | events.il-central-1.amazonaws.com<br>events.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | events.mx-central-1.amazonaws.com<br>events.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | events.me-south-1.amazonaws.com<br>events.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | events.me-central-1.amazonaws.com<br>events.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | events.sa-east-1.amazonaws.com<br>events.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | events.us-gov-east-1.amazonaws.com<br>events.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | events.us-gov-west-1.amazonaws.com<br>events.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Api destinations | Each supported Region: 3,000 | Yes | The maximum number of API destinations per account per Region. |
| Connections | Each supported Region: 3,000 | Yes | The maximum number of connections per account per Region. |
| CreateEndpoint throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of requests per second for CreateEndpoint API. Additional requests are throttled. |
| DeleteEndpoint throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of requests per second for DeleteEndpoint API. Additional requests are throttled. |
| Endpoints | Each supported Region: 100 | Yes | The maximum number of endpoints per account per Region. |
| Event bus policy size | Each supported Region: 10,240 | Yes | Maximum policy size, in characters. This policy size increases each time you grant access to another account. You can see your current policy and its size by using the DescribeEventBus API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Event buses | Each supported Region: 100 | Yes | Maximum event buses per account. |
| Event pattern size | Each supported Region: 2,048 | Yes | Maximum size of an event pattern, in characters. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Invocations throttle limit in transactions per second | us-east-1: 18,750 per second<br><br>us-east-2: 4,500 per second<br><br>us-west-1: 2,250 per second<br><br>us-west-2: 18,750 per second<br><br>ap-east-1: 1,100 per second<br><br>ap-northeast-1: 2,250 per second<br><br>ap-northeast-2: 1,100 per second<br><br>ap-south-1: 1,100 per second<br><br>ap-southeast-1: 2,250 per second<br><br>ap-southeast-2: 2,250 per second<br><br>ca-central-1: 1,100 per second<br><br>eu-central-1: 4,500 per second | Yes | An invocation is an event matching a rule and being sent on to the rules targets. After the limit is reached, the invocations are throttled; that is, they still happen but they are delayed. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-north-1: 1,100 per second<br><br>eu-south-2: 4,500 per second<br><br>eu-west-1: 18,750 per second<br><br>eu-west-2: 2,250 per second<br><br>eu-west-3: 1,100 per second<br><br>sa-east-1: 1,100 per second<br><br>Each of the other supported Regions: 750 per second | | |
| Number of rules | af-south-1: 100<br><br>eu-south-1: 100<br><br>Each of the other supported Regions: 300 | Yes | Maximum number of rules an account can have per event bus |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| PutEvents throttle limit in transactions per second | us-east-1: 10,000 per second<br><br>us-east-2: 2,400 per second<br><br>us-west-1: 1,200 per second<br><br>us-west-2: 10,000 per second<br><br>ap-east-1: 600 per second<br><br>ap-northeast-1: 1,200 per second<br><br>ap-northeast-2: 600 per second<br><br>ap-south-1: 600 per second<br><br>ap-southeast-1: 1,200 per second<br><br>ap-southeast-2: 1,200 per second<br><br>ca-central-1: 600 per second<br><br>eu-central-1: 2,400 per second | Yes | Maximum number of requests per second for PutEvents API. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-north-1: 600 per second<br><br>eu-south-2: 2,400 per second<br><br>eu-west-1: 10,000 per second<br><br>eu-west-2: 1,200 per second<br><br>eu-west-3: 600 per second<br><br>sa-east-1: 600 per second<br><br>Each of the other supported Regions: 400 per second | | |
| Rate of invocations per API destination | Each supported Region: 300 per second | Yes | The maximum number of invocations per second to send to each API destination endpoint per account per Region. Once the quota is met, future invocations to that API endpoint are throttled. The invocations will still occur, but are delayed. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Targets per rule | Each supported Region: 5 | No | Maximum number of targets that can be associated with a rule |
| Throttle limit in transactions per second for control plane APIs | Each supported Region: 50 per second | No | Maximum number of requests per second for EventBridge control plane API operations. Additional requests are throttled. |
| UpdateEndpoint throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of requests per second for UpdateEndpoint API. Additional requests are throttled. |

For more information, see CloudWatch Events quotas in the *Amazon CloudWatch Events User Guide*.

## Internet Monitor endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | internetmonitor.us-east-2.amazonaws.com | HTTPS | |
| | | internetmonitor-fips.us-east-2.api.aws | HTTPS | |
| | | internetmonitor-fips.us-east-2.amazonaws.com | HTTPS | |
| | | internetmonitor.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | internetmonitor.us-east-1.amazonaws.com | HTTPS | |
| | | internetmonitor-fips.us-east-1.api.aws | HTTPS | |
| | | internetmonitor-fips.us-east-1.amazonaws.com | HTTPS | |
| | | internetmonitor.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | internetmonitor.us-west-1.amazonaws.com | HTTPS | |
| | | internetmonitor-fips.us-west-1.api.aws | HTTPS | |
| | | internetmonitor-fips.us-west-1.amazonaws.com | HTTPS | |
| | | internetmonitor.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | internetmonitor.us-west-2.amazonaws.com | HTTPS | |
| | | internetmonitor-fips.us-west-2.api.aws | HTTPS | |
| | | internetmonitor-fips.us-west-2.amazonaws.com | HTTPS | |
| | | internetmonitor.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | internetmonitor.af-south-1.amazonaws.com<br>internetmonitor.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | internetmonitor.ap-east-1.amazonaws.com<br>internetmonitor.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | internetmonitor.ap-south-2.amazonaws.com<br>internetmonitor.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | internetmonitor.ap-southeast-3.amazonaws.com<br>internetmonitor.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | internetmonitor.ap-southeast-4.amazonaws.com<br>internetmonitor.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | internetmonitor.ap-south-1.amazonaws.com<br>internetmonitor.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | internetmonitor.ap-northeast-3.amazonaws.com<br>internetmonitor.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | internetmonitor.ap-northeast-2.amazonaws.com<br>internetmonitor.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | internetmonitor.ap-southeast-1.amazonaws.com<br>internetmonitor.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | internetmonitor.ap-southeast-2.amazonaws.com<br>internetmonitor.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | internetmonitor.ap-northeast-1.amazonaws.com<br>internetmonitor.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | internetmonitor.ca-central-1.amazonaws.com<br>internetmonitor-fips.ca-central-1.api.aws<br>internetmonitor-fips.ca-central-1.amazonaws.com<br>internetmonitor.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | internetmonitor.eu-central-1.amazonaws.com<br>internetmonitor.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | internetmonitor.eu-west-1.amazonaws.com<br>internetmonitor.eu-west-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | internetmonitor.eu-west-2.amazonaws.com | HTTPS |
| | | internetmonitor.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | internetmonitor.eu-south-1.amazonaws.com | HTTPS |
| | | internetmonitor.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | internetmonitor.eu-west-3.amazonaws.com | HTTPS |
| | | internetmonitor.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | internetmonitor.eu-south-2.amazonaws.com | HTTPS |
| | | internetmonitor.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | internetmonitor.eu-north-1.amazonaws.com | HTTPS |
| | | internetmonitor.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | internetmonitor.eu-central-2.amazonaws.com | HTTPS |
| | | internetmonitor.eu-central-2.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | internetmonitor.me-south-1.amazonaws.com | HTTPS |
| | | internetmonitor.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | internetmonitor.me-central-1.amazonaws.com | HTTPS |
| | | internetmonitor.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | internetmonitor.sa-east-1.amazonaws.com | HTTPS |
| | | internetmonitor.sa-east-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Days that health events are retained | Each supported Region: 400 | No | The number of days that AWS keeps information about a resolved Internet Monitor health event. |
| Monitors per account per AWS Region | Each supported Region: 50 | No | The maximum number of monitors in an account in one AWS Region. |
| Resources per monitor | Each supported Region: 50 | No | The maximum number of resources that a monitor can have. |

# Amazon CloudWatch Logs endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

> **Note**
>
> CloudWatch Logs also offers VPC Endpoints which support all of these endpoint types. For more information, see Using CloudWatch Logs with interface VPC endpoints in the *Amazon CloudWatch Logs User Guide*.

You can use the endpoints in the following tables to connect to APIs.

# Logs Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | logs.us-east-2.amazonaws.com | HTTPS | |
| | | logs.us-east-2.api.aws | HTTPS | |
| | | logs-fips.us-east-2.amazonaws.com | HTTPS | |
| | | logs-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | logs.us-east-1.amazonaws.com | HTTPS | |
| | | logs-fips.us-east-1.amazonaws.com | HTTPS | |
| | | logs.us-east-1.api.aws | HTTPS | |
| | | logs-fips.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | logs.us-west-1.amazonaws.com | HTTPS | |
| | | logs.us-west-1.api.aws | HTTPS | |
| | | logs-fips.us-west-1.api.aws | HTTPS | |
| | | logs-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | logs.us-west-2.amazonaws.com | HTTPS | |
| | | logs-fips.us-west-2.amazonaws.com | HTTPS | |
| | | logs.us-west-2.api.aws | HTTPS | |
| | | logs-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | logs.af-south-1.amazonaws.com | HTTPS | |
| | | logs.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | logs.ap-east-1.amazonaws.com<br>logs.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | logs.ap-south-2.amazonaws.com<br>logs.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | logs.ap-southeast-3.amazonaws.com<br>logs.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | logs.ap-southeast-5.amazonaws.com<br>logs.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | logs.ap-southeast-4.amazonaws.com<br>logs.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | logs.ap-south-1.amazonaws.com<br>logs.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | logs.ap-southeast-6.amazonaws.com<br>logs.ap-southeast-6.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | logs.ap-northeast-3.amazonaws.com<br>logs.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | logs.ap-northeast-2.amazonaws.com<br>logs.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | logs.ap-southeast-1.amazonaws.com<br>logs.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | logs.ap-southeast-2.amazonaws.com<br>logs.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | logs.ap-east-2.amazonaws.com<br>logs.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | logs.ap-southeast-7.amazonaws.com<br>logs.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | logs.ap-northeast-1.amazonaws.com<br>logs.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | logs.ca-central-1.amazonaws.com | HTTPS | |
| | | logs.ca-central-1.api.aws | HTTPS | |
| | | logs-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | logs-fips.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | logs.ca-west-1.amazonaws.com | HTTPS | |
| | | logs.ca-west-1.api.aws | HTTPS | |
| | | logs-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | logs-fips.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | logs.eu-central-1.amazonaws.com | HTTPS | |
| | | logs.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | logs.eu-west-1.amazonaws.com | HTTPS | |
| | | logs.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | logs.eu-west-2.amazonaws.com | HTTPS | |
| | | logs.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | logs.eu-south-1.amazonaws.com | HTTPS | |
| | | logs.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | logs.eu-west-3.amazonaws.com | HTTPS | |
| | | logs.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | logs.eu-south-2.amazonaws.com | HTTPS | |
| | | logs.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | logs.eu-north-1.amazonaws.com<br>logs.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | logs.eu-central-2.amazonaws.com<br>logs.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | logs.il-central-1.amazonaws.com<br>logs.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | logs.mx-central-1.amazonaws.com<br>logs.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | logs.me-south-1.amazonaws.com<br>logs.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | logs.me-central-1.amazonaws.com<br>logs.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | logs.sa-east-1.amazonaws.com<br>logs.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | logs.us-gov-east-1.amazonaws.com<br>logs.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | logs.us-gov-west-1.amazonaws.com<br>logs.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Active export task | Each supported Region: 1 | No | The number of active (running or pending) export tasks per account |
| AssociateKmsKey throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of associate-kms-key calls per second per account/per region |
| AssociateSourceToS3TableIntegration throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of AssociateSourceToS3TableIntegration calls per second per account/per region |
| AssociateSourceToS3TableIntegration throttle limit in transactions per second in a burst | Each supported Region: 10 | No | The maximum number of AssociateSourceToS3TableIntegration calls per second per account/per region in a burst |
| Batch size | Each supported Region: 1 Megabytes | No | The maximum batch size in MB of a put-log-events request |
| CancelExportTask throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of cancel-export-task calls per second per account/per region |
| CancelImportTask throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of CancelImportTask calls |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | per second per account/per region |
| CancelImportTask throttle limit in transactions per second in a burst | Each supported Region: 1 | No | The maximum number of CancelImportTask calls per second per account/per region in a burst |
| CreateExportTask throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of create-export-task calls per second per account/per region |
| CreateImportTask throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of CreateImportTask calls per second per account/per region |
| CreateImportTask throttle limit in transactions per second in a burst | Each supported Region: 1 | No | The maximum number of CreateImportTask calls per second per account/per region in a burst |
| CreateLogGroup throttle limit in transactions per second | Each supported Region: 10 per second | Yes | The maximum number of create-log-group calls per second per account/per region |
| CreateLogStream throttle limit in transactions per second | Each supported Region: 50 per second | Yes | The maximum number of create-log-stream calls per second per account/per region |
| Data archiving | Each supported Region: 5 Gigabytes | No | The maximum size in GB of free data archiving |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteDataProtectionPolicy throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of delete-data-protection-policy calls per second per account/per region |
| DeleteDestination throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of delete-destination calls per second per account/per region |
| DeleteLogGroup throttle limit in transactions per second | Each supported Region: 10 per second | Yes | The maximum number of delete-log-group calls per second per account/per region |
| DeleteLogStream throttle limit in transactions per second | Each supported Region: 15 per second | Yes | The maximum number of delete-log-stream calls per second per account/per region |
| DeleteMetricFilter throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of delete-metric-filter calls per second per account/per region |
| DeleteRetentionPolicy throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of delete-retention-policy calls per second per account/per region |
| DeleteSubscriptionFilter throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of delete-subscription-filter calls per second per account/per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DeleteTransformer throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of delete-transformer calls per second per account/per region |
| DescribeDestinations throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of describe-destinations calls per second per account/per region |
| DescribeExportTasks throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of describe-export-tasks calls per second per account/per region |
| DescribeImportTaskBatches throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeImportTaskBatches calls per second per account/per region |
| DescribeImportTaskBatches throttle limit in transactions per second in a burst | Each supported Region: 10 | No | The maximum number of DescribeImportTaskBatches calls per second per account/per region in a burst |
| DescribeImportTasks throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeImportTasks calls per second per account/per region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DescribeImportTasks throttle limit in transactions per second in a burst | Each supported Region: 10 | No | The maximum number of DescribeImportTasks calls per second per account/per region in a burst |
| DescribeLogGroups throttle limit in transactions per second | Each supported Region: 10 per second | Yes | The maximum number of describe-log-groups calls per second per account/per region |
| DescribeLogStreams throttle limit in transactions per second | Each supported Region: 25 per second | Yes | The maximum number of describe-log-streams calls per second per account/per region |
| DescribeMetricFilters throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of describe-metric-filters calls per second per account/per region |
| DescribeSubscriptionFilters throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of describe-subscription-filters calls per second per account/per region |
| DisassociateSourceFromS3TableIntegration throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of DisassociateSourceFromS3TableIntegration calls per second per account/per region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DisassociateSourceFromS3TableIntegration throttle limit in transactions per second in a burst | Each supported Region: 10 | No | The maximum number of DisassociateSourceFromS3TableIntegration calls per second per account/per region in a burst |
| Event size | Each supported Region: 1,024 Kilobytes | No | The maximum log event size in KB |
| FilterLogEvents throttle limit in transactions per second | us-east-1: 25 per second<br><br>ap-northeast-3: 5 per second<br><br>ap-southeast-3: 5 per second<br><br>ca-west-1: 5 per second<br><br>eu-central-1: 5 per second<br><br>il-central-1: 5 per second<br><br>Each of the other supported Regions: 10 per second | No | The maximum number of filter-log-events calls per second per account/per region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetDataProtectionPolicy throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of get-data-protection-policy calls per second per account/per region |
| GetLogEvents throttle limit in transactions per second | us-west-2: 10 per second<br><br>ap-northeast-3: 10 per second<br><br>ap-southeast-3: 10 per second<br><br>ca-west-1: 10 per second<br><br>eu-central-1: 10 per second<br><br>eu-west-1: 10 per second<br><br>eu-west-3: 30 per second<br><br>il-central-1: 10 per second<br><br>Each of the other supported Regions: 25 per second | No | The maximum number of get-log-events calls per second per account/per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| GetQueryResults throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of get-query-results calls per second per account/per region |
| GetTransformer throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of get-transformer calls per second per account/per region |
| ListSourcesForS3TableIntegration throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of ListSourcesForS3TableIntegration calls per second per account/per region |
| ListSourcesForS3TableIntegration throttle limit in transactions per second in a burst | Each supported Region: 10 | No | The maximum number of ListSourcesForS3TableIntegration calls per second per account/per region in a burst |
| ListTagsForResource throttle limit in transactions per second | Each supported Region: 30 per second | No | The maximum number of list-tags-for-resource calls per second per account/per region |
| ListTagsLogGroup throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of list-tags-log-group calls per second per account/per region |
| Live Tail concurrent sessions limit | Each supported Region: 15 | Yes | The maximum number of concurrent active Live Tail sessions per account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Log groups | Each supported Region: 1,000,000 | Yes | The maximum number of log groups an account can have |
| Log groups scanned per Live Tail session | Each supported Region: 10 | No | The maximum number of log groups that can be scanned per Live Tail session |
| Metrics filters per log group | Each supported Region: 100 | No | The number of metric filters per log group |
| PutBearerTokenAuthentication throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of put-bearer-token-authentication calls per second per account/per region |
| PutBearerTokenAuthentication throttle limit in transactions per second in a burst | Each supported Region: 10 | No | The maximum number of put-bearer-token-authentication calls per second per account/per region in a burst |
| PutDataProtectionPolicy throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of put-data-protection-policy calls per second per account/per region |
| PutDestination throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of put-destination calls per second per account/per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| PutDestinationPolicy throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of put-destination-policy calls per second per account/per region |
| PutLogEvents throttle limit in transactions per second | Each supported Region: 5,000 per second | Yes | The maximum number of put-log-events calls per second per account/per region |
| PutMetricFilter throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of put-metric-filter calls per second per account/per region |
| PutRetentionPolicy throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of put-retention-policy calls per second per account/per region |
| PutSubscriptionFilter throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of put-subscription-filter calls per second per account/per region |
| PutTransformer throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of put-transformer calls per second per account/per region |
| Resource policies | Each supported Region: 10 | No | The maximum number of resource policies per account/per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| StartLiveTail throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of start-live-tail calls per second per account/per region. This limit applies independently to both console and API |
| StartQuery throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of start-query calls per second per account/per region |
| Subscription filters per log group | Each supported Region: 5 | No | The number of subscription filters per log group |
| TagLogGroup throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of tag-log-group calls per second per account/per region |
| TagResource throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of tag-resource calls per second per account/per region |
| TestMetricFilter throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of test-metric-filter calls per second per account/per region |
| TestTransformer throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of test-transformer calls per second per account/per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| UntagLogGroup throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of untag-log-group calls per second per account/per region |
| UntagResource throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of untag-resource calls per second per account/per region |
| logs anomaly detectors | Each supported Region: 500 | Yes | The maximum number of active log anomaly detectors |

For more information, see CloudWatch Logs quotas in the *Amazon CloudWatch Logs User Guide*.

# Network Synthetic Monitor

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | networkmonitor.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | networkmonitor-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | networkmonitor.us-east-1.amazonaws.com<br>networkmonitor-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | networkmonitor.us-west-1.amazonaws.com<br>networkmonitor-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | networkmonitor.us-west-2.amazonaws.com<br>networkmonitor-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | networkmonitor.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | networkmonitor.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | networkmonitor.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | networkmonitor.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | | | |
| Asia Pacific (Melbourne) | ap-southeast-4 | networkmonitor.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | networkmonitor.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | networkmonitor.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | networkmonitor.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | networkmonitor.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | networkmonitor.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | networkmonitor.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | networkmonitor.ca-central-1.amazonaws.com<br>networkmonitor-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | | | |
| Europe (Frankfurt) | eu-central-1 | networkmonitor.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | networkmonitor.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | networkmonitor.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | networkmonitor.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | networkmonitor.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | networkmonitor.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | networkmonitor.eu-north-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | networkmonitor.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | networkmonitor.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | | | |
| Middle East (Bahrain) | me-south-1 | networkmonitor.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | networkmonitor.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | networkmonitor.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of monitors per account per AWS region | Each supported Region: 100 | Yes | The maximum number of monitors in an account in one AWS Region. |
| Number of probes per monitor | Each supported Region: 24 | Yes | The maximum number of probes that a monitor can have. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of probes per subnet for each monitor | Each supported Region: 4 | Yes | The maximum number of probes that a subnet in a monitor can have. |

# CloudWatch Observability Access Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | oam.us-east-2.amazonaws.com | HTTPS | |
| | | oam-fips.us-east-2.api.aws | HTTPS | |
| | | oam-fips.us-east-2.amazonaws.com | HTTPS | |
| | | oam.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | oam.us-east-1.amazonaws.com | HTTPS | |
| | | oam-fips.us-east-1.api.aws | HTTPS | |
| | | oam-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | oam.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | oam.us-west-1.amazonaws.com | HTTPS |
| | | oam-fips.us-west-1.api.aws | HTTPS |
| | | oam-fips.us-west-1.amazonaws.com | HTTPS |
| | | oam.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | oam.us-west-2.amazonaws.com | HTTPS |
| | | oam-fips.us-west-2.api.aws | HTTPS |
| | | oam-fips.us-west-2.amazonaws.com | HTTPS |
| | | oam.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | oam.af-south-1.amazonaws.com | HTTPS |
| | | oam.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | oam.ap-east-1.amazonaws.com | HTTPS |
| | | oam.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | oam.ap-south-2.amazonaws.com | HTTPS |
| | | oam.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | oam.ap-southeast-3.amazonaws.com | HTTPS |
| | | oam.ap-southeast-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | oam.ap-southeast-5.amazonaws.com<br>oam.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | oam.ap-southeast-4.amazonaws.com<br>oam.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | oam.ap-south-1.amazonaws.com<br>oam.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | oam.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | oam.ap-northeast-3.amazonaws.com<br>oam.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | oam.ap-northeast-2.amazonaws.com<br>oam.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | oam.ap-southeast-1.amazonaws.com<br>oam.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | oam.ap-southeast-2.amazonaws.com<br>oam.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | oam.ap-east-2.amazonaws.com | HTTPS |
| | | oam.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | oam.ap-southeast-7.amazonaws.com | HTTPS |
| | | oam.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | oam.ap-northeast-1.amazonaws.com | HTTPS |
| | | oam.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | oam.ca-central-1.amazonaws.com | HTTPS |
| | | oam-fips.ca-central-1.api.aws | HTTPS |
| | | oam-fips.ca-central-1.amazonaws.com | HTTPS |
| | | oam.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | oam.ca-west-1.amazonaws.com | HTTPS |
| | | oam-fips.ca-west-1.api.aws | HTTPS |
| | | oam-fips.ca-west-1.amazonaws.com | HTTPS |
| | | oam.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | oam.eu-central-1.amazonaws.com | HTTPS |
| | | oam.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | oam.eu-west-1.amazonaws.com | HTTPS |
| | | oam.eu-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | oam.eu-west-2.amazonaws.com | HTTPS |
| | | oam.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | oam.eu-south-1.amazonaws.com | HTTPS |
| | | oam.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | oam.eu-west-3.amazonaws.com | HTTPS |
| | | oam.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | oam.eu-south-2.amazonaws.com | HTTPS |
| | | oam.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | oam.eu-north-1.amazonaws.com | HTTPS |
| | | oam.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | oam.eu-central-2.amazonaws.com | HTTPS |
| | | oam.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | oam.il-central-1.amazonaws.com | HTTPS |
| | | oam.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | oam.mx-central-1.amazonaws.com | HTTPS |
| | | oam.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | oam.me-south-1.amazonaws.com | HTTPS |
| | | oam.me-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | oam.me-central-1.amazonaws.com | HTTPS |
| | | oam.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | oam.sa-east-1.amazonaws.com | HTTPS |
| | | oam.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | oam.us-gov-east-1.amazonaws.com | HTTPS |
| | | oam-fips.us-gov-east-1.api.aws | HTTPS |
| | | oam-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | oam.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | oam.us-gov-west-1.amazonaws.com | HTTPS |
| | | oam-fips.us-gov-west-1.api.aws | HTTPS |
| | | oam-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | oam.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Links per sink | Each supported Region: 100,000 | No | Maximum number of links that can be attached to a sink |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of links | Each supported Region: 5 | No | Maximum number of links in your account |
| Number of sinks | Each supported Region: 1 | No | Maximum number of sinks in your account |
| Rate of CreateLink requests | Each supported Region: 10 per second | No | Maximum number of CreateLink requests you can make per second, in this account in the current region |
| Rate of CreateSink requests | Each supported Region: 10 per second | No | Maximum number of CreateSink requests you can make per second, in this account in the current region |
| Rate of DeleteLink requests | Each supported Region: 10 per second | No | Maximum number of DeleteLink requests you can make per second, in this account in the current region |
| Rate of DeleteSink requests | Each supported Region: 10 per second | No | Maximum number of DeleteSink requests you can make per second, in this account in the current region |
| Rate of GetLink requests | Each supported Region: 10 per second | No | Maximum number of GetLink requests you can make per second, in this account in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetSink requests | Each supported Region: 10 per second | No | Maximum number of GetSink requests you can make per second, in this account in the current region |
| Rate of GetSinkPolicy requests | Each supported Region: 10 per second | No | Maximum number of GetSinkPolicy requests you can make per second, in this account in the current region |
| Rate of ListAttachedLinks requests | Each supported Region: 10 per second | No | Maximum number of ListAttachedLinks requests you can make per second, in this account in the current region |
| Rate of ListLinks requests | Each supported Region: 10 per second | No | Maximum number of ListLinks requests you can make per second, in this account in the current region |
| Rate of ListSinks requests | Each supported Region: 10 per second | No | Maximum number of ListSinks requests you can make per second, in this account in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListTagsForResource requests | Each supported Region: 10 per second | No | Maximum number of ListTagsForResource requests you can make per second, in this account in the current region |
| Rate of PutSinkPolicy requests | Each supported Region: 1 per second | No | Maximum number of PutSinkPolicy requests you can make per second, in this account in the current region |
| Rate of TagResource requests | Each supported Region: 10 per second | No | Maximum number of TagResource requests you can make per second, in this account in the current region |
| Rate of UntagResource requests | Each supported Region: 10 per second | No | Maximum number of UntagResource requests you can make per second, in this account in the current region |
| Rate of UpdateLink requests | Each supported Region: 10 per second | No | Maximum number of UpdateLink requests you can make per second, in this account in the current region |

# Amazon CloudWatch Observability Admin endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | observabilityadmin.us-east-2.amazonaws.com | | |
| US East (N. Virginia) | us-east-1 | observabilityadmin.us-east-1.amazonaws.com | | |
| US West (N. California) | us-west-1 | observabilityadmin.us-west-1.amazonaws.com | | |
| US West (Oregon) | us-west-2 | observabilityadmin.us-west-2.amazonaws.com | | |
| Africa (Cape Town) | af-south-1 | observabilityadmin.af-south-1.amazonaws.com | | |
| Asia Pacific | ap-east-1 | observabilityadmin.ap-east-1.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | observabilityadmin.ap-south-2.amazonaws.com | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | observabilityadmin.ap-southeast-3.amazonaws.com | | |
| Asia Pacific (Malaysia) | ap-southeast-5 | observabilityadmin.ap-southeast-5.amazonaws.com | | |
| Asia Pacific (Melbourne) | ap-southeast-4 | observabilityadmin.ap-southeast-4.amazonaws.com | | |
| Asia Pacific (Mumbai) | ap-south-1 | observabilityadmin.ap-south-1.amazonaws.com | | |
| Asia Pacific (New Zealand) | ap-southeast-6 | observabilityadmin.ap-southeast-6.amazonaws.com | | |
| Asia Pacific (Osaka) | ap-northeast-3 | observabilityadmin.ap-northeast-3.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | observabilityadmin.ap-northeast-2.amazonaws.com | | |
| Asia Pacific (Singapore) | ap-southeast-1 | observabilityadmin.ap-southeast-1.amazonaws.com | | |
| Asia Pacific (Sydney) | ap-southeast-2 | observabilityadmin.ap-southeast-2.amazonaws.com | | |
| Asia Pacific (Taipei) | ap-east-2 | observabilityadmin.ap-east-2.amazonaws.com | | |
| Asia Pacific (Thailand) | ap-southeast-7 | observabilityadmin.ap-southeast-7.amazonaws.com | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | observabilityadmin.ap-northeast-1.amazonaws.com | | |
| Canada (Central) | ca-central-1 | observabilityadmin.ca-central-1.amazonaws.com | | |
| Canada West (Calgary) | ca-west-1 | observabilityadmin.ca-west-1.amazonaws.com | | |
| Europe (Frankfurt) | eu-central-1 | observabilityadmin.eu-central-1.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | observabilityadmin.eu-west-1.amazonaws.com | | |
| Europe (London) | eu-west-2 | observabilityadmin.eu-west-2.amazonaws.com | | |
| Europe (Milan) | eu-south-1 | observabilityadmin.eu-south-1.amazonaws.com | | |
| Europe (Paris) | eu-west-3 | observabilityadmin.eu-west-3.amazonaws.com | | |
| Europe (Spain) | eu-south-2 | observabilityadmin.eu-south-2.amazonaws.com | | |
| Europe (Stockholm) | eu-north-1 | observabilityadmin.eu-north-1.amazonaws.com | | |
| Europe (Zurich) | eu-central-2 | observabilityadmin.eu-central-2.amazonaws.com | | |
| Israel (Tel Aviv) | il-central-1 | observabilityadmin.il-central-1.amazonaws.com | | |
| Mexico (Central) | mx-central-1 | observabilityadmin.mx-central-1.amazonaws.com | | |
| South America (São Paulo) | sa-east-1 | observabilityadmin.sa-east-1.amazonaws.com | | |
| AWS GovCloud (US-East) | us-gov-east-1 | observabilityadmin.us-gov-east-1.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | observabilityadmin.us-gov-west-1.amazonaws.com | | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CreateCentralizationRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of create-centralization-rule-for-organization calls per second per account/per region |
| CreateTelemetryPipeline throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of CreateTelemetryPipeline calls per second per account and per Region |
| CreateTelemetryRule throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of create-telemetry-rule calls per second per account/per region |
| CreateTelemetryRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of create-telemetry-rule-for-organization calls per second per account/per region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteCentralizationRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of delete-centralization-rule-for-organization calls per second per account/per region |
| DeleteTelemetryPipeline throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of DeleteTelemetryPipeline calls per second per account and per Region |
| DeleteTelemetryRule throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of delete-telemetry-rule calls per second per account/per region |
| DeleteTelemetryRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of delete-telemetry-rule-for-organization calls per second per account/per region |
| GetCentralizationRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of get-centralization-rule-for-organization calls per second per account/per region |
| GetTelemetryEnrichmentStatus throttle limit in transactions per second | Each supported Region: 5 per second | Yes | The maximum number of GetTelemetryEnrichmentStatus calls per second per account and per region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| GetTelemetryPipeline throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of GetTelemetryPipeline calls per second per account and per Region |
| GetTelemetryRule throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of get-telemetry-rule calls per second per account/per region |
| GetTelemetryRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of get-telemetry-rule-for-organization calls per second per account/per region |
| ListCentralizationRulesForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of list-centralization-rules-for-organization calls per second per account/per region |
| ListTagsForResource throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of list-tags-for-resource calls per second per account/per region |
| ListTelemetryPipelines throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of ListTelemetryPipelines calls per second per account and per Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListTelemetryRules throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of list-telemetry-rules calls per second per account/per region |
| ListTelemetryRulesForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of list-telemetry-rules-for-organization calls per second per account/per region |
| Logs Centralization throughput per destination region | Each supported Region: 2.5 Gigabytes per second | No | The maximum organization-level throughput for centralization of logs per destination region |
| Organization Centralization Rules | Each supported Region: 50 | Yes | The maximum number of organization centralization rules that you can create for an organization |
| Organization Telemetry Rules | Each supported Region: 50 | Yes | The maximum number of organization telemetry rules that you can create for an organization |
| Pipelines for managed OTel metrics experience | Each supported Region: 300 | Yes | The maximum number of pipelines for managed OTel metrics experience allowed per account and per Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Pipelines with a CloudWatch Logs source | Each supported Region: 300 | Yes | The maximum number of pipelines with a CloudWatch Logs source allowed per account per Region |
| Pipelines with a third-party integration source | Each supported Region: 30 | Yes | The maximum number of pipelines with a third-party integration source allowed per account per Region |
| StartTelemetryEnrichment throttle limit in transactions per second | Each supported Region: 1 per second | Yes | The maximum number of StartTelemetryEnrichment calls per second per account and per region |
| StopTelemetryEnrichment throttle limit in transactions per second | Each supported Region: 1 per second | Yes | The maximum number of StopTelemetryEnrichment calls per second per account and per region |
| TagResource throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of tag-resource calls per second per account/per region |
| Telemetry Rules | Each supported Region: 20 | Yes | The maximum number of telemetry rules that you can create for an account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| TestTelemetryPipeline throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of TestTelemetryPipeline calls per second per account and per Region |
| UntagResource throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of untag-resource calls per second per account/per region |
| UpdateCentralizationRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of update-centralization-rule-for-organization calls per second per account/per region |
| UpdateTelemetryPipeline throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of UpdateTelemetryPipeline calls per second per account and per Region |
| UpdateTelemetryRule throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of update-telemetry-rule calls per second per account/per region |
| UpdateTelemetryRuleForOrganization throttle limit in transactions per second | Each supported Region: 5 | No | The maximum number of update-telemetry-rule-for-organization calls per second per account/per region |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| ValidateTelemetryPipelineConfiguration throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of ValidateTelemetryPipelineConfiguration calls per second per account and per Region |

# Amazon CloudWatch RUM endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | rum.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | rum.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | rum.us-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | rum.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | rum.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | rum.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | rum.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rum.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rum.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rum.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rum.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | rum.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | rum.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | rum.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | rum.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | rum.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | rum.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | rum.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | rum.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | rum.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | rum.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | rum.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | rum.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | rum.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | rum.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | rum.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | rum.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | rum.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | rum.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | rum.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | rum.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | rum.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | rum.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | rum.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RUM AppMonitors | Each supported Region: 20 | Yes | The maximum number of RUM AppMonitors that you can have in this AWS account. |
| RUM Events per second per AWS Account | Each supported Region: 50 | Yes | The maximum number of RUM Events per second that RUM will ingest for this AWS account. |

For more information, see CloudWatch RUM quotas in the *Amazon CloudWatch User Guide*.

# Amazon CloudWatch Synthetics endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | synthetics.us-east-2.amazonaws.com | HTTPS |
| | | synthetics-fips.us-east-2.api.aws | HTTPS |
| | | synthetics-fips.us-east-2.amazonaws.com | HTTPS |
| | | synthetics.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | synthetics.us-east-1.amazonaws.com | HTTPS |
| | | synthetics-fips.us-east-1.api.aws | HTTPS |
| | | synthetics-fips.us-east-1.amazonaws.com | HTTPS |
| | | synthetics.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | synthetics.us-west-1.amazonaws.com | HTTPS |
| | | synthetics-fips.us-west-1.api.aws | HTTPS |
| | | synthetics-fips.us-west-1.amazonaws.com | HTTPS |
| | | synthetics.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | synthetics.us-west-2.amazonaws.com | HTTPS |
| | | synthetics-fips.us-west-2.api.aws | HTTPS |
| | | synthetics-fips.us-west-2.amazonaws.com | HTTPS |
| | | synthetics.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | synthetics.af-south-1.amazonaws.com | HTTPS |
| | | synthetics.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | synthetics.ap-east-1.amazonaws.com | HTTPS |
| | | synthetics.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | synthetics.ap-south-2.amazonaws.com | HTTPS |
| | | synthetics.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | synthetics.ap-southeast-3.amazonaws.com | HTTPS |
| | | synthetics.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | synthetics.ap-southeast-5.amazonaws.com | HTTPS |
| | | synthetics.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | synthetics.ap-southeast-4.amazonaws.com | HTTPS |
| | | synthetics.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | synthetics.ap-south-1.amazonaws.com<br>synthetics.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | synthetics.ap-southeast-6.amazonaws.com<br>synthetics.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | synthetics.ap-northeast-3.amazonaws.com<br>synthetics.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | synthetics.ap-northeast-2.amazonaws.com<br>synthetics.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | synthetics.ap-southeast-1.amazonaws.com<br>synthetics.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | synthetics.ap-southeast-2.amazonaws.com<br>synthetics.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | synthetics.ap-east-2.amazonaws.com<br>synthetics.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | synthetics.ap-southeast-7.amazonaws.com<br>synthetics.ap-southeast-7.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | synthetics.ap-northeast-1.amazonaws.com | HTTPS |
| | | synthetics.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | synthetics.ca-central-1.amazonaws.com | HTTPS |
| | | synthetics-fips.ca-central-1.api.aws | HTTPS |
| | | synthetics-fips.ca-central-1.amazonaws.com | HTTPS |
| | | synthetics.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | synthetics.ca-west-1.amazonaws.com | HTTPS |
| | | synthetics-fips.ca-west-1.api.aws | HTTPS |
| | | synthetics-fips.ca-west-1.amazonaws.com | HTTPS |
| | | synthetics.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | synthetics.eu-central-1.amazonaws.com | HTTPS |
| | | synthetics.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | synthetics.eu-west-1.amazonaws.com | HTTPS |
| | | synthetics.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | synthetics.eu-west-2.amazonaws.com | HTTPS |
| | | synthetics.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | synthetics.eu-south-1.amazonaws.com | HTTPS |
| | | synthetics.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | synthetics.eu-west-3.amazonaws.com | HTTPS |
| | | synthetics.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | synthetics.eu-south-2.amazonaws.com | HTTPS | |
| | | synthetics.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | synthetics.eu-north-1.amazonaws.com | HTTPS | |
| | | synthetics.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | synthetics.eu-central-2.amazonaws.com | HTTPS | |
| | | synthetics.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | synthetics.il-central-1.amazonaws.com | HTTPS | |
| | | synthetics.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | synthetics.mx-central-1.amazonaws.com | HTTPS | |
| | | synthetics.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | synthetics.me-south-1.amazonaws.com | HTTPS | |
| | | synthetics.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | synthetics.me-central-1.amazonaws.com | HTTPS | |
| | | synthetics.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | synthetics.sa-east-1.amazonaws.com | HTTPS | |
| | | synthetics.sa-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | synthetics.us-gov-east-1.amazonaws.com | HTTPS |
| | | synthetics-fips.us-gov-east-1.api.aws | HTTPS |
| | | synthetics-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | synthetics.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | synthetics.us-gov-west-1.amazonaws.com | HTTPS |
| | | synthetics-fips.us-gov-west-1.api.aws | HTTPS |
| | | synthetics-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | synthetics.us-gov-west-1.api.aws | HTTPS |

### Service quotas

| Name | Default | Adjustable |
|---|---|---|
| Number of canaries | 200 per Region per account. | Yes |

For more information, see CloudWatch service quotas in the *Amazon CloudWatch User Guide*.

## AWS CodeArtifact endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | codeartifact.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | codeartifact.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | codeartifact.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | codeartifact.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codeartifact.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codeartifact.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codeartifact.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codeartifact.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | codeartifact.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codeartifact.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | codeartifact.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | codeartifact.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | codeartifact.eu-north-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Asset file size | Each supported Region: 5 Gigabytes | Yes | The maximum file size per asset. |
| Assets per package version | Each supported Region: 350 | No | The maximum number of assets per package version. |
| CopyPackageVersions requests per second | Each supported Region: 5 | Yes | The maximum number of calls that can be made to CopyPackageVersions per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Direct upstreams per repository | Each supported Region: 10 | No | The maximum number of direct upstream repositories per repository. |
| Domains per AWS account | Each supported Region: 10 | Yes | The maximum number of domains that can be created per AWS account. |
| GetAuthorizationToken requests per second | Each supported Region: 40 | Yes | The maximum number of authorization tokens retrieved per second. |
| GetPackageVersionAsset requests per second | Each supported Region: 50 | Yes | The maximum number of calls that can be made to GetPackageVersionAsset per second. |
| ListPackageVersionAssets requests per second | Each supported Region: 200 | Yes | The maximum number of calls that can be made to ListPackageVersionAssets per second. |
| ListPackageVersions requests per second | Each supported Region: 200 | Yes | The maximum number of calls that can be made to ListPackageVersions per second. |
| ListPackages requests per second | Each supported Region: 200 | Yes | The maximum number of calls that can be made to ListPackages per second. |
| PublishPackageVersion requests per second | Each supported Region: 10 | Yes | The maximum number of calls that can be made to PublishPackageVersion per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Read requests per second from a single AWS account | Each supported Region: 800 | Yes | The maximum number of read requests from one AWS account per second. |
| Repositories per domain | Each supported Region: 1,000 | Yes | The maximum number of repositories that can be created per domain. |
| Requests per second using a single authentication token | Each supported Region: 1,200 | No | The maximum number of requests per second using a single authentication token. |
| Requests without authentication token per IP address | Each supported Region: 600 | No | The maximum number of requests per second without an authentication token from a single IP address. |
| Upstream repositories searched | Each supported Region: 25 | No | The maximum number of upstream repositories searched when resolving a package. |
| Write requests per second from a single AWS account | Each supported Region: 100 | Yes | The maximum number of write requests from one AWS account per second. |

# AWS CodeBuild endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | codebuild.us-east-2.amazonaws.com<br>codebuild-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | codebuild.us-east-1.amazonaws.com<br>codebuild-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | codebuild.us-west-1.amazonaws.com<br>codebuild-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | codebuild.us-west-2.amazonaws.com<br>codebuild-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | codebuild.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | codebuild.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-south-2 | codebuild.ap-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | codebuild.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | codebuild.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | codebuild.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | codebuild.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | codebuild.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | codebuild.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codebuild.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | codebuild.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codebuild.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | codebuild.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codebuild.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | codebuild.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codebuild.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codebuild.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codebuild.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | codebuild.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | codebuild.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | codebuild.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | codebuild.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | codebuild.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | codebuild.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | codebuild.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | codebuild.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | codebuild.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | codebuild.us-gov-east-1.amazonaws.com<br>codebuild-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | codebuild.us-gov-west-1.amazonaws.com<br>codebuild-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Associated tags per project | Each supported Region: 50 | No | Maximum number of tags you can associate with a build project |
| Build projects | Each supported Region: 5,000 | Yes | Maximum number of build projects |
| Build timeout in minutes | Each supported Region: 2,160 | No | Maximum build timeout in minutes |
| Concurrent request for information about builds | Each supported Region: 100 | No | Maximum number of builds you can request information about at any one time using the AWS CLI or an AWS SDK. |
| Concurrent requests for information on build projects | Each supported Region: 100 | No | Maximum number of build projects you can request information about at any one time using the AWS CLI or an AWS SDK. |
| Concurrently running builds for ARM Lambda/10GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM Lambda/10GB environment |
| Concurrently running builds for ARM Lambda/1GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | Lambda/1GB environment |
| Concurrently running builds for ARM Lambda/2GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM Lambda/2GB environment |
| Concurrently running builds for ARM Lambda/4GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM Lambda/4GB environment |
| Concurrently running builds for ARM Lambda/8GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM Lambda/8GB environment |
| Concurrently running builds for ARM/2XLarge environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM/2XLarge environment |
| Concurrently running builds for ARM/Large environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM/Large environment |
| Concurrently running builds for ARM/Medium environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM/Medium environment |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrently running builds for ARM/Small environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM/Small environment |
| Concurrently running builds for ARM/XLarge environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for ARM/XLarge environment |
| Concurrently running builds for Linux GPU Large environment | Each supported Region: 0 | Yes | Maximum number of concurrently running builds for Linux GPU/Large environment |
| Concurrently running builds for Linux GPU Small environment | Each supported Region: 0 | Yes | Maximum number of concurrently running builds for Linux GPU/Small environment |
| Concurrently running builds for Linux Lambda/10GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux Lambda/10GB environment |
| Concurrently running builds for Linux Lambda/1GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux Lambda/1GB environment |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrently running builds for Linux Lambda/2GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux Lambda/2GB environment |
| Concurrently running builds for Linux Lambda/4GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux Lambda/4GB environment |
| Concurrently running builds for Linux Lambda/8GB environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux Lambda/8GB environment |
| Concurrently running builds for Linux/2XLarge environment | Each supported Region: 0 | Yes | Maximum number of concurrently running builds for Linux/2XLarge environment |
| Concurrently running builds for Linux/Large environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux/Large environment |
| Concurrently running builds for Linux/Medium environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux/Medium environment |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrently running builds for Linux/Small environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux/Small environment |
| Concurrently running builds for Linux/XLarge environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Linux/XLarge environment |
| Concurrently running builds for Windows Server 2019/Large environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows Server 2019/Large environment |
| Concurrently running builds for Windows Server 2019/Medium environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows Server 2019/Medium environment |
| Concurrently running builds for Windows Server 2022/2XLarge environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows Server 2022/2XLarge environment |
| Concurrently running builds for Windows Server 2022/Large environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows Server 2022/Large environment |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Concurrently running builds for Windows Server 2022/Medium environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows Server 2022/Medium environment |
| Concurrently running builds for Windows Server 2022/XLarge environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows Server 2022/XLarge environment |
| Concurrently running builds for Windows/Large environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows/Large environment |
| Concurrently running builds for Windows/Medium environment | Each supported Region: 1 | Yes | Maximum number of concurrently running builds for Windows/Medium environment |
| Minimum period for build timeout in minutes | Each supported Region: 5 | No | Minimum build timeout in minutes |
| Security groups under VPC configuration | Each supported Region: 5 | No | Security groups available for VPC configuration |
| Subnets under VPC configuration | Each supported Region: 16 | No | Subnets available for VPC configuration |

**Note**

Internal metrics will determine the default quotas for concurrent running builds.

For more information, see Quotas for CodeBuild in the *AWS CodeBuild User Guide*.

# AWS CodeCommit endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | codecommit.us-east-2.amazonaws.com | HTTPS |
| | | codecommit-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | codecommit.us-east-1.amazonaws.com | HTTPS |
| | | codecommit-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | codecommit.us-west-1.amazonaws.com | HTTPS |
| | | codecommit-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | codecommit.us-west-2.amazonaws.com | HTTPS |
| | | codecommit-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | codecommit.af-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | codecommit.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | codecommit.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | codecommit.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | codecommit.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | codecommit.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codecommit.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codecommit.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codecommit.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | codecommit.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | codecommit.ca-central-1.amazonaws.com | HTTPS |
| | | codecommit-fips.ca-central-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | codecommit.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | codecommit.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | codecommit.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | codecommit.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | codecommit.eu-west-3.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | codecommit.eu-north-1.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | codecommit.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | codecommit.me-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | codecommit.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | codecommit.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | codecommit.us-gov-east-1.amazonaws.com<br>codecommit-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | codecommit.us-gov-west-1.amazonaws.com<br>codecommit-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

For information about Git connection endpoints, including SSH and HTTPS information, see Regions and Git Connection Endpoints for CodeCommit.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Allowed repositories | Each supported Region: 5,000 | Yes | The maximum number of repositories that you can create in this account. |

For more information, see Quotas in CodeCommit in the *AWS CodeCommit User Guide*.

# CodeConnections endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region Name | Region | Endpoint | Protocol |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | codeconnections.us-east-2.amazonaws.com | HTTPS | |
| | | codeconnections-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | codeconnections.us-east-1.amazonaws.com | HTTPS | |
| | | codeconnections-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | codeconnections.us-west-1.amazonaws.com | HTTPS | |
| | | codconnections-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | codeconnections.us-west-2.amazonaws.com | HTTPS | |
| | | codeconnections-fips.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region Name | Region | Endpoint | Protocol |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | codeconnections.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codeconnections.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codeconnections.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codeconnections.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codeconnections.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | codeconnections.ca-central-1.amazonaws.com<br>codeconnections-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codeconnections.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codeconnections.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codeconnections.eu-west-2.amazonaws.com | HTTPS | |

| Region Name | Region Name | Region | Endpoint | Protocol |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | codeconnections.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | codeconnections.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | codeconnections.eu-north-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | codeconnections.sa-east-1.amazonaws.com | HTTPS | |
| GovCloud (US-East) | us-gov-east-1 | codeconnections.us-gov-east-1.amazonaws.com<br>codeconnections-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |

### Service quotas

For a list of quotas, see Quotas for connections in the *Developer Tools Console User Guide*.

## AWS CodeDeploy endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | codedeploy.us-east-2.amazonaws.com<br>codedeploy-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | codedeploy.us-east-1.amazonaws.com<br>codedeploy-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | codedeploy.us-west-1.amazonaws.com<br>codedeploy-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | codedeploy.us-west-2.amazonaws.com<br>codedeploy-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | codedeploy.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | codedeploy.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | codedeploy.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | codedeploy.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | codedeploy.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | codedeploy.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | codedeploy.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | codedeploy.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | codedeploy.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codedeploy.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codedeploy.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codedeploy.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | codedeploy.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | codedeploy.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codedeploy.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | codedeploy.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | codedeploy.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codedeploy.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codedeploy.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codedeploy.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | codedeploy.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | codedeploy.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | codedeploy.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | codedeploy.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | codedeploy.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | codedeploy.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | codedeploy.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | codedeploy.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | codedeploy.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | codedeploy.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | codedeploy.us-gov-east-1.amazonaws.com | HTTPS |
| | | codedeploy-fips.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-West) | us-gov-west-1 | codedeploy.us-gov-west-1.amazonaws.com<br><br>codedeploy-fips.us-gov-west-1.amazonaws.com | HTTPS<br><br>HTTPS |

## CodeDeploy agent endpoints

> **Note**
>
> These endpoints support IPv4 only.

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | codedeploy-commands.us-east-2.amazonaws.com<br><br>codedeploy-commands-fips.us-east-2.amazonaws.com<br><br>codedeploy-commands-secure.us-east-2.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS |
| US East (N. Virginia) | us-east-1 | codedeploy-commands.us-east-1.amazonaws.com<br><br>codedeploy-commands-fips.us- | HTTPS<br><br>HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| | | east-1.amazonaws.com<br><br>codedeploy-commands-secure.us-east-1.amazonaws.com | |
| US West (N. California) | us-west-1 | codedeploy-commands.us-west-1.amazonaws.com<br><br>codedeploy-commands-fips.us-west-1.amazonaws.com<br><br>codedeploy-commands-secure.us-west-1.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS |
| US West (Oregon) | us-west-2 | codedeploy-commands.us-west-2.amazonaws.com<br><br>codedeploy-commands-fips.us-west-2.amazonaws.com<br><br>codedeploy-commands-secure.us-west-2.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Africa (Cape Town) | af-south-1 | codedeploy-commands.af-south-1.amazonaws.com<br>codedeploy-commands-secure.af-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | codedeploy-commands.ap-east-1.amazonaws.com<br>codedeploy-commands-secure.ap-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | codedeploy-commands.ap-east-2.amazonaws.com<br>codedeploy-commands-secure.ap-east-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | codedeploy-commands.ap-south-2.amazonaws.com<br>codedeploy-commands-secure.ap-south-2.amazonaws.com | HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Jakarta) | ap-southeast-3 | codedeploy-commands.ap-southeast-3.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-3.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | codedeploy-commands.ap-southeast-4.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-4.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | codedeploy-commands.ap-southeast-5.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-5.amazonaws.com | HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (New Zealand) | ap-southeast-6 | codedeploy-commands.ap-southeast-6.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-6.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | codedeploy-commands.ap-southeast-7.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-7.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | codedeploy-commands.ap-south-1.amazonaws.com<br><br>codedeploy-commands-secure.ap-south-1.amazonaws.com | HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Osaka) | ap-northeast-3 | codedeploy-commands.ap-northeast-3.amazonaws.com<br><br>codedeploy-commands-secure.ap-northeast-3.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | codedeploy-commands.ap-northeast-2.amazonaws.com<br><br>codedeploy-commands-secure.ap-northeast-2.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | codedeploy-commands.ap-southeast-1.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-1.amazonaws.com | HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Sydney) | ap-southeast-2 | codedeploy-commands.ap-southeast-2.amazonaws.com<br><br>codedeploy-commands-secure.ap-southeast-2.amazonaws.com | HTTPS<br><br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | codedeploy-commands.ap-northeast-1.amazonaws.com<br><br>codedeploy-commands-secure.ap-northeast-1.amazonaws.com | HTTPS<br><br>HTTPS |
| Canada (Central) | ca-central-1 | codedeploy-commands.ca-central-1.amazonaws.com<br><br>codedeploy-commands-secure.ca-central-1.amazonaws.com | HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Canada West (Calgary) | ca-west-1 | codedeploy-commands.ca-west-1.amazonaws.com<br><br>codedeploy-commands-secure.ca-west-1.amazonaws.com | HTTPS<br><br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | codedeploy-commands.eu-central-1.amazonaws.com<br><br>codedeploy-commands-secure.eu-central-1.amazonaws.com | HTTPS<br><br>HTTPS |
| Europe (Ireland) | eu-west-1 | codedeploy-commands.eu-west-1.amazonaws.com<br><br>codedeploy-commands-secure.eu-west-1.amazonaws.com | HTTPS<br><br>HTTPS |
| Europe (London) | eu-west-2 | codedeploy-commands.eu-west-2.amazonaws.com<br><br>codedeploy-commands-secure.eu-west-2.amazonaws.com | HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Milan) | eu-south-1 | codedeploy-commands.eu-south-1.amazonaws.com<br>codedeploy-commands-secure.eu-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | codedeploy-commands.eu-west-3.amazonaws.com<br>codedeploy-commands-secure.eu-west-3.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | codedeploy-commands.eu-south-2.amazonaws.com<br>codedeploy-commands-secure.eu-south-2.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | codedeploy-commands.eu-north-1.amazonaws.com<br>codedeploy-commands-secure.eu-north-1.amazonaws.com | HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Zurich) | eu-central-2 | codedeploy-commands.eu-central-2.amazonaws.com<br>codedeploy-commands-secure.eu-central-2.amazonaws.com | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | codedeploy-commands.il-central-1.amazonaws.com<br>codedeploy-commands-secure.il-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | codedeploy-commands.me-south-1.amazonaws.com<br>codedeploy-commands-secure.me-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | codedeploy-commands.me-central-1.amazonaws.com<br>codedeploy-commands-secure.me-central-1.amazonaws.com | HTTPS<br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Mexico (Central) | mx-central-1 | codedeploy-commands.mx-central-1.amazonaws.com<br><br>codedeploy-commands-secure.mx-central-1.amazonaws.com | HTTPS<br><br>HTTPS |
| South America (São Paulo) | sa-east-1 | codedeploy-commands.sa-east-1.amazonaws.com<br><br>codedeploy-commands-secure.sa-east-1.amazonaws.com | HTTPS<br><br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | codedeploy-commands.us-gov-east-1.amazonaws.com<br><br>codedeploy-commands-fips.us-gov-east-1.amazonaws.com<br><br>codedeploy-commands-secure.us-gov-east-1.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-West) | us-gov-west-1 | codedeploy-commands.us-gov-west-1.amazonaws.com<br>codedeploy-commands-fips.us-gov-west-1.amazonaws.com<br>codedeploy-commands-secure.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| AWS Lambda deployment run in hours | Each supported Region: 50 | No | Maximum number of hours an AWS Lambda deployment can run (48 hours for the maximum time between the first and last traffic shift plus one hour for each of two possible lifecycle hooks) |
| Applications associated per account per region | Each supported Region: 1,000 | Yes | The maximum number of applications associated with an AWS account in a single region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Associated alarms per deployment group | Each supported Region: 50 | Yes | The maximum number of alarms associated with a deployment group |
| Auto Scaling groups in a deployment group | Each supported Region: 10 | Yes | Maximum number of Amazon EC2 Auto Scaling groups in a deployment group |
| Concurrent deployments per account | Each supported Region: 1,300 | Yes | Maximum number of concurrent deployments associated with an AWS account. Each deployment to a scaled-up Amazon EC2 instance in an Amazon EC2 Auto Scaling group counts as a single concurrent deployment for each application that the EC2 instance is associated with. |
| Concurrent deployments per deployment group | Each supported Region: 1 | No | Maximum number of concurrent deployments to a deployment group. This limit prevents concurrent deployments of the same application to the same deployment group. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom deployment configurations per account | Each supported Region: 50 | No | Maximum number of custom deployment configurations associated with an AWS account |
| Deployment groups associated with a single application | Each supported Region: 1,000 | Yes | Maximum number of deployment groups associated with a single application |
| EC2/On-Premises blue/green deployment run in hours | Each supported Region: 109 | No | Maximum number of hours an EC2/On-Premises blue/green deployment can run (48 hours for each of the above two limits plus one hour for each of 13 possible lifecycle events) |
| EC2/On-Premises in-place deployment run in hours | Each supported Region: 12 | No | Maximum number of hours an EC2/On-Premises in-place deployment can run |
| Event notification triggers in a deployment group | Each supported Region: 10 | Yes | Maximum number of event notification triggers in a deployment group |
| GitHub connection tokens per account | Each supported Region: 25 | No | Maximum number of GitHub connection tokens for a single AWS account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Hours between the completion of a deployment and the termination of the original instances during an EC2/On-Premises blue/green deployment | Each supported Region: 48 | No | Maximum number of hours between the completion of a deployment and the termination of the original instances during an EC2/On-Premises blue/green deployment |
| Hours between the deployment of a revision and when traffic shifts to the replacement instances during an EC2/On-Premises blue/green deployment | Each supported Region: 48 | No | Maximum number of hours between the deployment of a revision and when traffic shifts to the replacement instances during an EC2/On-Premises blue/green deployment |
| Instances count per deployment | us-east-1: 2,000<br>us-east-2: 2,000<br>us-west-2: 2,000<br>eu-central-1: 2,000<br>Each of the other supported Regions: 1,000 | Yes | Maximum number of instances in a single deployment |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Minutes a blue/green deployment can wait after a successful deployment before terminating instances from the original deployment | Each supported Region: 2,800 | No | Maximum number of minutes a blue/green deployment can wait after a successful deployment before terminating instances from the original deployment |
| Minutes between the first and last traffic shift during an AWS Lambda canary or linear deployment | Each supported Region: 2,880 | No | Maximum number of minutes between the first and last traffic shift during an AWS Lambda canary or linear deployment |
| Minutes until a deployment fails if a lifecycle event doesnt start | Each supported Region: 5 | No | Maximum number of minutes until a deployment fails if a lifecycle event doesnt start after (1) a deployment is triggered by using the console or the AWS CLI create-deployment command, or (2) the previous lifecycle event is completed. |
| Number of deployment groups that can be associated with an Amazon ECS service | Each supported Region: 1 | No | Maximum number of deployment groups that can be associated with an Amazon ECS service |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of instances that can be passed to the BatchGetOnPremisesInstances API action | Each supported Region: 100 | No | Maximum number of instances that can be passed to the BatchGetOnPremisesInstances API action |
| Number of instances used by concurrent deployments that are in progress per account | us-east-1: 3,000<br><br>Each of the other supported Regions: 1,000 | Yes | Maximum number of instances that can be used by concurrent deployments that are in progress and associated with one account |
| Number of listeners for a traffic route during an Amazon ECS deployment | Each supported Region: 1 | No | Maximum number of listeners for a traffic route during an Amazon ECS deployment |
| Seconds until a deployment lifecycle event fails if not completed | Each supported Region: 3,600 Seconds | No | Maximum number of seconds until a deployment lifecycle event fails if not completed |
| Size of deployment group name | Each supported Region: 100 | No | Maximum number of characters in a deployment group name |
| Size of tag key | Each supported Region: 128 | No | Maximum number of characters in a tag key |
| Size of tag value | Each supported Region: 256 | No | Maximum number of characters in a tag value |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Tags in a deployment group | Each supported Region: 10 | No | Maximum number of tags in a deployment group |
| Traffic that can be shifted in one increment during an AWS Lambda deployment | Each supported Region: 99 | No | Maximum percentage of traffic that can be shifted in one increment during an AWS Lambda deployment |

For more information, see Quotas in CodeDeploy in the *AWS CodeDeploy User Guide*.

# Amazon CodeGuru Profiler endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | codeguru-profiler.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | codeguru-profiler.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | codeguru-profiler.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codeguru-profiler.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codeguru-profiler.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codeguru-profiler.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codeguru-profiler.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codeguru-profiler.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codeguru-profiler.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | codeguru-profiler.eu-north-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Number of profiling groups per account and region. | Each supported Region: 500 | Yes | The maximum number of profiling groups per account, per region. |

# Amazon CodeGuru Reviewer endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | codeguru-reviewer.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | codeguru-reviewer.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | codeguru-reviewer.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | codeguru-reviewer.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codeguru-reviewer.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codeguru-reviewer.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codeguru-reviewer.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codeguru-reviewer.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codeguru-reviewer.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | codeguru-reviewer.eu-north-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Allowed Code Reviews | Each supported Region: 5,000 | Yes | The maximum number of code reviews per account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | per month in the current region. |

# Amazon CodeGuru Security endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | codeguru-security.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | codeguru-security.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | codeguru-security.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | codeguru-security.us-east-2.amazonaws.com | HTTPS |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Sydney) | ap-southeast-2 | codeguru-security.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | codeguru-security.ap-northeast-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | codeguru-security.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | codeguru-security.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | codeguru-security.eu-west-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | codeguru-security.eu-north-1.amazonaws.com | HTTPS |

# AWS CodePipeline endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | codepipeline.us-east-2.amazonaws.com<br>codepipeline-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | codepipeline.us-east-1.amazonaws.com<br>codepipeline-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | codepipeline.us-west-1.amazonaws.com<br>codepipeline-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | codepipeline.us-west-2.amazonaws.com<br>codepipeline-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | codepipeline.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | codepipeline.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | codepipeline.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | codepipeline.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | codepipeline.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | codepipeline.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | codepipeline.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | codepipeline.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | codepipeline.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codepipeline.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codepipeline.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codepipeline.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | codepipeline.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | codepipeline.ca-central-1.amazonaws.com<br>codepipeline-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codepipeline.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codepipeline.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codepipeline.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | codepipeline.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | codepipeline.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | codepipeline.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | codepipeline.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | codepipeline.eu-central-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | codepipeline.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | codepipeline.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | codepipeline.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | codepipeline.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | codepipeline.us-gov-east-1.amazonaws.com<br>codepipeline-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | codepipeline.us-gov-west-1.amazonaws.com<br>codepipeline-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AWS CloudFormation action timeout | Each supported Region: 3 | Yes | The length of time, in days, before an AWS CloudFormation |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | deployment action times out |
| AWS CodeBuild action timeout | Each supported Region: 36 | No | The length of time, in hours, before an AWS CodeBuild build action or test action times out |
| AWS CodeDeploy ECS (Blue/Green) action timeout | Each supported Region: 5 | Yes | The length of time, in days, before an AWS CodeDeploy ECS (Blue/Green) action times out |
| AWS CodeDeploy action timeout | Each supported Region: 5 | Yes | The length of time, in days, before an AWS CodeDeploy deployment action times out |
| AWS Lambda action timeout | Each supported Region: 24 | Yes | The length of time in hours before an AWS Lambda invoke action times out |
| Action configuration key length | Each supported Region: 50 | No | The maximum number of characters allowed in an action configuration key |
| Action configuration value length | Each supported Region: 1,000 | No | The maximum number of characters allowed in an action configuration value |
| Action timeout | Each supported Region: 1 | Yes | The length of time in hours before any other action times out |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Active executions per pipeline | Each supported Region: 50 | No | The maximum number of active executions allowed in a PARALLEL execution mode pipeline. |
| Amazon S3 deployment action timeout | Each supported Region: 90 | Yes | The length of time, in minutes, before an Amazon S3 deployment action times out |
| Approval action timeout | Each supported Region: 7 | No | The length of time, in days, before an approval action times out |
| Minimum actions | Each supported Region: 1 | No | The minimum number of actions allowed in a stage |
| Minimum stages per pipeline | Each supported Region: 2 | No | The minimum number of stages allowed in a pipeline. |
| Total AWS CodeCommit or GitHub source artifact size | Each supported Region: 1 Gigabytes | No | The maximum size (in GB) of artifacts in a source stage that uses AWS CodeCommit or GitHub repositories |
| Total Amazon S3 source artifact size | Each supported Region: 7 Gigabytes | No | The maximum size (in GB) of artifacts in a source stage that uses Amazon S3 artifact buckets |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Total JSON object size for Parameter Overrides | Each supported Region: 1 Kilobytes | No | The maximum size (in KB) of the JSON object that can be stored in the ParameterOverrides property |
| Total actions per pipeline | Each supported Region: 1,000 | No | The maximum number of actions allowed in a pipeline |
| Total actions per stage | Each supported Region: 100 | No | The maximum number of actions allowed in a stage |
| Total custom actions | Each supported Region: 50 | Yes | The maximum number of custom actions per region in an AWS account |
| Total image definitions JSON file size | Each supported Region: 100 Kilobytes | No | The maximum size (in KB) of the image definitions file JSON file used in pipelines that deploy Amazon ECS containers and images |
| Total input artifact size for AWS CloudFormation deployments | Each supported Region: 256 Megabytes | No | The maximum size (in MB) of input artifacts for AWS CloudFormation actions when deploying Lambda functions |
| Total parallel actions per stage | Each supported Region: 100 | No | The maximum number of parallel actions in a stage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Total period for execution history | Each supported Region: 12 | No | The maximum number of previous months for which pipeline execution history information can be viewed |
| Total pipelines | Each supported Region: 1,000 | Yes | The maximum number of pipelines per AWS Region in an AWS account. |
| Total pipelines with change detection set to periodically checking for source changes | Each supported Region: 300 | No | The maximum number of pipelines per region with change detection set to periodically checking for source changes |
| Total sequential actions per stage | Each supported Region: 100 | No | The maximum number of sequential actions in a stage |
| Total source artifact size for Amazon EBS deployments | Each supported Region: 512 Megabytes | No | The maximum size (in MB) of artifacts in a source stage for a pipeline that uses Amazon EBS to deploy applications |
| Total stages per pipeline | Each supported Region: 50 | No | The maximum number of stages allowed in a pipeline |
| Total webhooks | Each supported Region: 300 | Yes | The maximum number of webhooks per region in an AWS account |

For more information, see Quotas in CodePipeline in the *AWS CodePipeline User Guide*.

# AWS CodeStar Connections endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | codestar-connections.us-east-2.amazonaws.com<br><br>codestar-connections-fips.us-east-2.amazonaws.com | HTTPS<br><br>HTTPS |
| US East (N. Virginia) | us-east-1 | codestar-connections.us-east-1.amazonaws.com<br><br>codestar-connections-fips.us-east-1.amazonaws.com | HTTPS<br><br>HTTPS |
| US West (N. California) | us-west-1 | codestar-connections.us-west-1.amazonaws.com<br><br>codestar-connections-fips.us-west-1.amazonaws.com | HTTPS<br><br>HTTPS |
| US West (Oregon) | us-west-2 | codestar-connections.us-west-2.amazonaws.com | HTTPS<br><br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | codestar-connections-fips.us-west-2.amazonaws.com | | |
| Asia Pacific (Mumbai) | ap-south-1 | codestar-connections.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codestar-connections.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | codestar-connections.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codestar-connections.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codestar-connections.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | codestar-connections.ca-central-1.amazonaws.com<br>codestar-connections-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codestar-connections.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codestar-connections.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | codestar-connections.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | codestar-connections.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | codestar-connections.eu-west-3.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | codestar-connections.eu-north-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | codestar-connections.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | codestar-connections.us-gov-east-1.amazonaws.com<br><br>codestar-connections-fips.us-gov-east-1.amazonaws.com | HTTPS<br><br>HTTPS |

### Service quotas

For a list of quotas, see Quotas for connections in the *Developer Tools Console User Guide*.

## AWS CodeStar Notifications endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | codestar-notifications.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | codestar-notifications.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | codestar-notifications.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | codestar-notifications.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | codestar-notifications.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | codestar-notifications.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | codestar-notifications.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | codestar-notifications.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | codestar-notifications.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | codestar-notifications.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | codestar-notifications.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | codestar-notifications.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | codestar-notifications.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | codestar-notifications.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | codestar-notifications.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | codestar-notifications.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | codestar-notifications.me-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | codestar-notifications.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

For a list of quotas, see Quotas for notifications in the *Developer Tools console User Guide*.

# Amazon Cognito endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Amazon Cognito user pools

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cognito-idp.us-east-2.amazonaws.com | HTTPS | |
| | | cognito-idp-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | cognito-idp.us-east-1.amazonaws.com | HTTPS | |
| | | cognito-idp-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (N. California) | us-west-1 | cognito-idp.us-west-1.amazonaws.com<br>cognito-idp-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | cognito-idp.us-west-2.amazonaws.com<br>cognito-idp-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | cognito-idp.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cognito-idp.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | cognito-idp.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | cognito-idp.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | cognito-idp.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | cognito-idp.ap-southeast-4.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | cognito-idp.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | cognito-idp.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cognito-idp.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cognito-idp.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cognito-idp.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cognito-idp.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | cognito-idp.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | cognito-idp.ap-southeast-7.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | cognito-idp.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | cognito-idp.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | cognito-idp.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cognito-idp.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | cognito-idp.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cognito-idp.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | cognito-idp.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | cognito-idp.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | cognito-idp.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | cognito-idp.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | cognito-idp.eu-central-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | cognito-idp.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | cognito-idp.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | cognito-idp.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | cognito-idp.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | cognito-idp.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | cognito-idp.us-gov-east-1.amazonaws.com<br>cognito-idp-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | cognito-idp.us-gov-west-1.amazonaws.com<br>cognito-idp-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Amazon Cognito identity pools

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cognito-identity.us-east-2.amazonaws.com<br>cognito-identity-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | cognito-identity.us-east-1.amazonaws.com<br>cognito-identity-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | cognito-identity.us-west-1.amazonaws.com<br>cognito-identity-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | cognito-identity.us-west-2.amazonaws.com<br>cognito-identity-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | cognito-identity.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cognito-identity.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | cognito-identity.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | cognito-identity.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | cognito-identity.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | cognito-identity.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cognito-identity.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | cognito-identity.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | cognito-identity.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cognito-identity.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cognito-identity.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | cognito-identity.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | cognito-identity.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | cognito-identity.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cognito-identity.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | cognito-identity.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | cognito-identity.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cognito-identity.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | cognito-identity.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cognito-identity.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | cognito-identity.eu-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | cognito-identity.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | cognito-identity.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | cognito-identity.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | cognito-identity.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | cognito-identity.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | cognito-identity.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | cognito-identity.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | cognito-identity.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | cognito-identity.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | cognito-identity.us-gov-east-1.amazonaws.com<br><br>cognito-identity-fips.us-gov-east-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | cognito-identity.us-gov-west-1.amazonaws.com<br><br>cognito-identity-fips.us-gov-west-1.amazonaws.com | HTTPS<br><br>HTTPS | |

## Amazon Cognito Sync

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | cognito-sync.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | cognito-sync.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | cognito-sync.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | cognito-sync.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cognito-sync.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | cognito-sync.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cognito-sync.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cognito-sync.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cognito-sync.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | cognito-sync.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cognito-sync.eu-west-2.amazonaws.com | HTTPS | |

# Service quotas

## Amazon Cognito user pools

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Apps per user pool | Each supported Region: 1,000 | Yes | The maximum number of app clients per user pool. |
| Custom domains per account | Each supported Region: 4 | No | The maximum number of custom domains that |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | you can create in this account. |
| Groups per user | Each supported Region: 100 | No | The maximum number of groups per user that any individual user can be added to. |
| Groups per user pool | Each supported Region: 10,000 | No | The maximum number of groups per user pool. A group is a collection of users in a user pool. |
| Identity providers per user pool | Each supported Region: 300 | Yes | The maximum number of identity providers per user pool. |
| Rate of ClientAuthentication requests per account | Each supported Region: 150 per second | No | The maximum combined call rate (requests per second) for all API operations in the ClientAuthentication category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of LimitManagement requests | Each supported Region: 1 per second | No | The maximum combined call rate (requests per second) for all API operations in the LimitManagement category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserAccountRecovery requests | Each supported Region: 30 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserAccountRecovery category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserAuthentication requests | Each supported Region: 120 per second | Yes | The maximum combined call rate (requests per second) you can request for all API operations in the UserAuthentication category. Provision your rate limit within this quota from the Amazon Cognito console or APIs. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserCreation requests | Each supported Region: 50 per second | Yes | The maximum combined call rate (requests per second) you can request for all API operations in the UserCreation category. Provision your rate limit within this quota from the Amazon Cognito console or APIs. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserFederation requests | Each supported Region: 25 per second | Yes | The maximum combined call rate (requests per second) you can request for all API operations in the UserFederation category. Provision your rate limit within this quota from the Amazon Cognito console or APIs. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserList requests | Each supported Region: 30 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserList category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserPoolClientRead requests per account | Each supported Region: 15 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserPoolClientRead category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserPoolClientRead requests per user pool | Each supported Region: 5 per second | No | The maximum call rate (requests per second) for an operation in the UserPoolClientRead category per user pool. Any operation within this category could be called at this rate per user pool. You can find the list of included operations at https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserPoolClientUpdate requests per account | Each supported Region: 15 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserPoolClientUpdate category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserPoolClientUpdate requests per user pool | Each supported Region: 5 per second | No | The maximum call rate (requests per second) for an operation in the UserPoolClientUpdate category per user pool. Any operation within this category could be called at this rate per user pool. You can find the list of included operations at https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserPoolRead requests | Each supported Region: 15 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserPoolRead category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserPoolResourceRead requests per account | Each supported Region: 20 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserPoolResourceRead category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserPoolResourceRead requests per user pool | Each supported Region: 5 per second | No | The maximum call rate (requests per second) for an operation in the UserPoolResourceRead category per user pool. Any operation within this category could be called at this rate per user pool. You can find the list of included operations at https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserPoolResourceUpdate requests per account | Each supported Region: 15 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserPoolResourceUpdate category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserPoolResourceUpdate requests per user pool | Each supported Region: 5 per second | No | The maximum call rate (requests per second) for an operation in the UserPoolResourceUpdate category per user pool. Any operation within this category could be called at this rate per user pool. You can find the list of included operations at https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserPoolUpdate requests | Each supported Region: 15 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserPoolUpdate category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserRead requests | Each supported Region: 120 per second | Yes | The maximum combined call rate (requests per second) you can request for all API operations in the UserRead category. Provision your rate limit within this quota from the Amazon Cognito console or APIs. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserResourceRead requests | Each supported Region: 50 per second | Yes | The maximum combined call rate (requests per second) you can request for all API operations in the UserResourceRead category. Provision your rate limit within this quota from the Amazon Cognito console or APIs. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserResourceUpdate requests | Each supported Region: 25 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserResourceUpdate category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Rate of UserToken requests | Each supported Region: 120 per second | Yes | The maximum combined call rate (requests per second) you can request for all API operations in the UserToken category. Provision your rate limit within this quota from the Amazon Cognito console or APIs. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UserUpdate requests | Each supported Region: 25 per second | No | The maximum combined call rate (requests per second) for all API operations in the UserUpdate category. All operations within a category share the same quota. For the list of included operations, see https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html#category_operations. |
| Resource servers per user pool | Each supported Region: 25 | Yes | The maximum number of resource servers per user pool. A resource server is a server for access-protected resources. |
| Scopes per resource server | Each supported Region: 100 | No | The maximum number of scopes per resource server. A scope is a level of access (such as read or write access) that an app can request to a resource. |
| User import jobs per user pool | Each supported Region: 1,000 | Yes | The maximum number of user import jobs per user pool. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| User pools per account | Each supported Region: 1,000 | Yes | The maximum number of user pools that you can create in this account per region. |

For more information, see Quotas in Amazon Cognito in the *Amazon Cognito Developer Guide*.

## Amazon Cognito identity pools

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Identity pool name size | Each supported Region: 128 Bytes | No | The maximum size of an identity pool name in bytes. |
| Identity pools per account | Each supported Region: 1,000 | Yes | The maximum number of identity pools per account. |
| List API call results | Each supported Region: 60 | No | The maximum number of results from a list or lookup API call. |
| Login provider name size | Each supported Region: 2,048 Bytes | No | The maximum size of a login provider name in bytes. |
| Rate of GetCredentialsForIdentity requests | Each supported Region: 200 | Yes | The maximum number of GetCredentialsForIdentity requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetId requests | Each supported Region: 25 | Yes | The maximum number of GetId requests that you can make per second. |
| Rate of GetOpenIdToken requests | Each supported Region: 200 | Yes | The maximum number of GetOpenIdToken requests that you can make per second. |
| Rate of GetOpenIdTokenForDeveloperIdentity requests | Each supported Region: 50 | Yes | The maximum number of GetOpenIdTokenForDeveloperIdentity requests that you can make per second. |
| Rate of ListIdentities requests | Each supported Region: 5 | Yes | The maximum number of ListIdentities requests that you can make per second. |
| Rate of ListTagsForResource requests | Each supported Region: 10 | Yes | The maximum number of ListTagsForResource requests that you can make per second. |
| Rate of TagResource requests | Each supported Region: 5 | Yes | The maximum number of TagResource requests that you can make per second. |
| Rate of UntagResource requests | Each supported Region: 5 | Yes | The maximum number of UntagResource requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Role-based access control rules | Each supported Region: 25 | No | The maximum number of rules for role-based access control (RBAC) |
| User pool providers per identity pool | Each supported Region: 50 | Yes | The maximum number of Amazon Cognito user pool providers per identity pool. |

For more information, see Quotas in Amazon Cognito in the *Amazon Cognito Developer Guide*.

## Amazon Cognito Sync

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Bulk publish wait time | Each supported Region: 24 | No | The maximum wait time for a bulk publish after a successful request in hours. |
| Dataset name size | Each supported Region: 128 Bytes | No | The maximum size of a dataset name in bytes. |
| Dataset size | Each supported Region: 1 Megabytes | Yes | The maximum size of a dataset in megabytes. |
| Datasets per identity | Each supported Region: 20 | Yes | The maximum number of datasets per identity. |
| Records per dataset | Each supported Region: 1,024 | Yes | The maximum number of records per dataset. |

For more information, see Quotas in Amazon Cognito in the *Amazon Cognito Developer Guide*.

# Amazon Comprehend endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | comprehend.us-east-2.amazonaws.com | HTTPS |
| | | comprehend.us-east-2.api.aws | HTTPS |
| | | comprehend-fips.us-east-2.amazonaws.com | HTTPS |
| | | comprehend-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | comprehend.us-east-1.amazonaws.com | HTTPS |
| | | comprehend-fips.us-east-1.amazonaws.com | HTTPS |
| | | comprehend-fips.us-east-1.api.aws | HTTPS |
| | | comprehend.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | comprehend.us-west-2.amazonaws.com | HTTPS |
| | | comprehend-fips.us-west-2.amazonaws.com | HTTPS |
| | | comprehend.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | comprehend-fips.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | comprehend.ap-south-1.amazonaws.com | HTTPS |
| | | comprehend.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | comprehend.ap-northeast-2.amazonaws.com | HTTPS |
| | | comprehend.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | comprehend.ap-southeast-1.amazonaws.com | HTTPS |
| | | comprehend.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | comprehend.ap-southeast-2.amazonaws.com | HTTPS |
| | | comprehend.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | comprehend.ap-northeast-1.amazonaws.com | HTTPS |
| | | comprehend.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | comprehend.ca-central-1.amazonaws.com | HTTPS |
| | | comprehend.ca-central-1.api.aws | HTTPS |
| | | comprehend-fips.ca-central-1.amazonaws.com | HTTPS |
| | | comprehend-fips.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | comprehend.eu-central-1.amazonaws.com | HTTPS |
| | | comprehend.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | comprehend.eu-west-1.amazonaws.com | HTTPS |
| | | comprehend.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | comprehend.eu-west-2.amazonaws.com | HTTPS |
| | | comprehend.eu-west-2.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | comprehend.us-gov-west-1.amazonaws.com | HTTPS |
| | | comprehend-fips.us-gov-west-1.api.aws | HTTPS |
| | | comprehend-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | comprehend.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| BatchDetectDominantLanguage throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of BatchDetectDominantLanguage requests allowed per account per second, in the current Region |
| BatchDetectEntities throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of BatchDetectEntities requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BatchDetectKeyPhrases throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of BatchDetectKeyPhrases requests allowed per account per second, in the current Region |
| BatchDetectSentiment throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of BatchDetectSentiment requests allowed per account per second, in the current Region |
| BatchDetectSyntax throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of BatchDetectSyntax requests allowed per account per second, in the current Region |
| BatchDetectTargetedSentiment throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of BatchDetectTargetedSentiment requests allowed per account per second, in the current Region |
| ContainsPiiEntities throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of ContainsPiiEntities requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CreateDocumentClassifier max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API CreateDocumentClassifier allowed per account, in the current Region |
| CreateDocumentClassifier throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of CreateDocumentClassifier requests allowed per account per second, in the current Region |
| CreateEntityRecognizer max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API CreateEntityRecognizer allowed per account, in the current Region |
| CreateEntityRecognizer throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of CreateEntityRecognizer requests allowed per account per second, in the current Region |
| CreateFlywheel max active flywheels | Each supported Region: 20 | Yes | The maximum number of flywheels that can be active concurrently per account, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CreateFlywheel max concurrent | Each supported Region: 10 | Yes | The maximum number of flywheels that can be created concurrently per account, in the current region |
| Datasets max concurrent creates | Each supported Region: 10 | Yes | The maximum number of datasets in creating state per account, in the current Region |
| DeleteDocumentClassifier throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of DeleteDocumentClassifier requests allowed per account per second, in the current Region |
| DeleteEntityRecognizer throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of DeleteEntityRecognizer requests allowed per account per second, in the current Region |
| DescribeDocumentClassificationJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeDocumentClassificationJob requests allowed per account per second, in the current Region |
| DescribeDocumentClassifier throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeDocumentClassifier requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DescribeDominantLanguageDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeDominantLanguageDetectionJob requests allowed per account per second, in the current Region |
| DescribeEntitiesDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeEntitiesDetectionJob requests allowed per account per second, in the current Region |
| DescribeEntityRecognizer throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeEntityRecognizer requests allowed per account per second, in the current Region |
| DescribeEventsDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeEventsDetectionJob requests allowed per account per second, in the current Region |
| DescribeKeyPhrasesDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeKeyPhrasesDetectionJob requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DescribePiiEntitiesDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribePiiEntitiesDetectionJob requests allowed per account per second, in the current Region |
| DescribeSentimentDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeSentimentDetectionJob requests allowed per account per second, in the current Region |
| DescribeTargetedSentimentDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeTargetedSentimentDetectionJob requests allowed per account per second, in the current Region |
| DescribeTopicsDetectionJob throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of DescribeTopicsDetectionJob requests allowed per account per second, in the current Region |
| DetectDominantLanguage throttle limit in transactions per second | us-east-1: 200<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of DetectDominantLanguage requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DetectEntities throttle limit in transactions per second | us-east-1: 100<br><br>Each of the other supported Regions: 25 | Yes | The maximum number of DetectEntities requests allowed per account per second, in the current Region |
| DetectKeyPhrases throttle limit in transactions per second | us-east-1: 50<br><br>Each of the other supported Regions: 25 | Yes | The maximum number of DetectKeyPhrases requests allowed per account per second, in the current Region |
| DetectPiiEntities throttle limit in transactions per second | us-east-1: 200<br><br>us-west-2: 100<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of DetectPiiEntities requests allowed per account per second, in the current Region |
| DetectSentiment throttle limit in transactions per second | us-east-1: 200<br><br>us-west-2: 100<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of DetectSentiment requests allowed per account per second, in the current Region |
| DetectSyntax throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of DetectSyntax requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DetectTargetedSentiment throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of DetectTargetedSentiment requests allowed per account per second, in the current Region |
| DetectToxicContent throttle limit in transactions per second | Each supported Region: 25 | Yes | The maximum number of DetectToxicContent requests allowed per account per second, in the current Region |
| Endpoints max active endpoints | Each supported Region: 20 | Yes | The maximum number of active endpoints allowed per account in the current Region |
| Endpoints max inference units per account | Each supported Region: 200 | Yes | The maximum number of inference units allowed per account in the current Region |
| Endpoints max inference units per endpoint | Each supported Region: 50 | Yes | The maximum number of inference units allowed per endpoint in the current Region |
| ListDocumentClassificationJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListDocumentClassificationJobs requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListDocumentClassifiers throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListDocumentClassifiers requests allowed per account per second, in the current Region |
| ListDominantLanguageDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListDominantLanguageDetectionJobs requests allowed per account per second, in the current Region |
| ListEntitiesDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListEntitiesDetectionJobs requests allowed per account per second, in the current Region |
| ListEntityRecognizers throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListEntityRecognizers requests allowed per account per second, in the current Region |
| ListEventsDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListEventsDetectionJobs requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ListKeyPhrasesDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListKeyPhrasesDetectionJobs requests allowed per account per second, in the current Region |
| ListPiiEntitiesDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListPiiEntitiesDetectionJobs requests allowed per account per second, in the current Region |
| ListSentimentDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListSentimentDetectionJobs requests allowed per account per second, in the current Region |
| ListTagsForResource throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListTagsForResource requests allowed per account per second, in the current Region |
| ListTargetedSentimentDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListTargetedSentimentDetectionJobs requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ListTopicsDetectionJobs throttle limit in transactions per second | Each supported Region: 10 | No | The maximum number of ListTopicsDetectionJobs requests allowed per account per second, in the current Region |
| MaxTestDatasets per flywheel | Each supported Region: 50 | Yes | The maximum number of test datasets per flywheel |
| MaxTrainDatasets per flywheel | Each supported Region: 50 | Yes | The maximum number of training datasets per flywheel |
| StartDocumentClassificationJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartDocumentClassificationJob allowed per account, in the current Region |
| StartDocumentClassificationJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartDocumentClassificationJob requests allowed per account per second, in the current Region |
| StartDominantLanguageDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartDominantLanguageDetectionJob allowed per account, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| StartDominantLanguageDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartDominantLanguageDetectionJob requests allowed per account per second, in the current Region |
| StartEntitiesDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartEntitiesDetectionJob allowed per account, in the current Region |
| StartEntitiesDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartEntitiesDetectionJob requests allowed per account per second, in the current Region |
| StartEventsDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartEventsDetectionJob allowed per account, in the current Region |
| StartEventsDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartEventsDetectionJob requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StartFlywheelIteration max concurrent flywheel iterations | Each supported Region: 10 | Yes | The maximum number of flywheel iterations that you can run concurrently per account, in the current Region |
| StartKeyPhrasesDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartKeyPhrasesDetectionJob allowed per account, in the current Region |
| StartKeyPhrasesDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartKeyPhrasesDetectionJob requests allowed per account per second, in the current Region |
| StartPiiEntitiesDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartPiiEntitiesDetectionJob allowed per account, in the current Region |
| StartPiiEntitiesDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartPiiEntitiesDetectionJob requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StartSentimentDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartSentimentDetectionJob allowed per account, in the current Region |
| StartSentimentDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartSentimentDetectionJob requests allowed per account per second, in the current Region |
| StartTargetedSentimentDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartTargetedSentimentDetectionJob allowed per account, in the current Region |
| StartTargetedSentimentDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartTargetedSentimentDetectionJob requests allowed per account per second, in the current Region |
| StartTopicsDetectionJob max active jobs | Each supported Region: 10 | Yes | The maximum number of active jobs for API StartTopicsDetectionJob allowed per account, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| StartTopicsDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StartTopicsDetectionJob requests allowed per account per second, in the current Region |
| StopDominantLanguageDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopDominantLanguageDetectionJob requests allowed per account per second, in the current Region |
| StopEntitiesDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopEntitiesDetectionJob requests allowed per account per second, in the current Region |
| StopEventsDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopEventsDetectionJob requests allowed per account per second, in the current Region |
| StopKeyPhrasesDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopKeyPhrasesDetectionJob requests allowed per account per second, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| StopPiiEntitiesDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopPiiEntitiesDetectionJob requests allowed per account per second, in the current Region |
| StopSentimentDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopSentimentDetectionJob requests allowed per account per second, in the current Region |
| StopTargetedSentimentDetectionJob throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopTargetedSentimentDetectionJob requests allowed per account per second, in the current Region |
| StopTrainingDocumentClassifier throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopTrainingDocumentClassifier requests allowed per account per second, in the current Region |
| StopTrainingEntityRecognizer throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of StopTrainingEntityRecognizer requests allowed per account per second, in the current Region |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| TagResource throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of TagResource requests allowed per account per second, in the current Region |
| UntagResource throttle limit in transactions per second | Each supported Region: 1 | No | The maximum number of UntagResource requests allowed per account per second, in the current Region |

For more information, see Guidelines and Quotas in the *Amazon Comprehend Developer Guide*.

# Amazon Comprehend Medical endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | comprehendmedical.us-east-2.amazonaws.com | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | comprehendmedical-fips.us-east-2.amazonaws.com | | |
| US East (N. Virginia) | us-east-1 | comprehendmedical.us-east-1.amazonaws.com<br>comprehendmedical-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | comprehendmedical.us-west-2.amazonaws.com<br>comprehendmedical-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | comprehendmedical.ap-southeast-2.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | comprehendmedical.ca-central-1.amazonaws.com<br>comprehendmedical-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | comprehendmedical.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | comprehendmedical.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | comprehendmedical.us-gov-west-1.amazonaws.com<br>comprehendmedical-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Characters per second (CPS) for the DetectEntities operation | Each supported Region: 40,000 | Yes | The maximum characters per second (CPS) for the DetectEntities operation. |
| Characters per second (CPS) for the DetectEntities-v2 operation | Each supported Region: 40,000 | Yes | The maximum characters per second (CPS) for the DetectEntities-v2 operation. |
| Characters per second (CPS) for the DetectPHI operation | Each supported Region: 40,000 | Yes | The maximum characters per second (CPS) for the DetectPHI operation. |
| Characters per second (CPS) for the InferICD10CM operation | Each supported Region: 40,000 | Yes | The maximum characters per second (CPS) for the InferICD10CM operation. |
| Characters per second (CPS) for the InferRxNorm operation | Each supported Region: 40,000 | Yes | The maximum characters per second (CPS) for the InferRxNorm operation. |
| Characters per second (CPS) for the InferSNOMEDCT operation | Each supported Region: 5,000 | Yes | The maximum characters per second (CPS) for the InferSNOMEDCT operation. |
| Maximum batch job size (sum total for all the files submitted in a batch job). | Each supported Region: 1 Gigabytes | No | The maximum batch job size (sum total for all the files submitted in a batch job). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum document size (UTF-8 characters) for InferSNOMEDCT | Each supported Region: 5 Kilobytes | No | The maximum document size (UTF-8 characters) for InferSNOMEDCT. |
| Maximum document size (UTF-8 characters) for the DetectEntities operation | Each supported Region: 20 Kilobytes | No | The maximum document size (UTF-8 characters) for the DetectEntities operation. |
| Maximum document size (UTF-8 characters) for the DetectEntities-v2 operation | Each supported Region: 20 Kilobytes | No | The maximum document size (UTF-8 characters) for the DetectEntities-v2 operation. |
| Maximum document size (UTF-8 characters) for the DetectPHI operation | Each supported Region: 20 Kilobytes | No | The maximum document size (UTF-8 characters) for the DetectPHI operation. |
| Maximum document size (UTF-8 characters) for the InferICD10CM operation | Each supported Region: 10 Kilobytes | No | The maximum document size (UTF-8 characters) for the InferICD10CM operation. |
| Maximum document size (UTF-8 characters) for the InferRxNorm operation | Each supported Region: 10 Kilobytes | No | The maximum document size (UTF-8 characters) for the InferRxNorm operation. |
| Maximum individual file size for InferSNOMEDCT batch jobs | Each supported Region: 5 Kilobytes | No | The maximum individual file size for InferSNOMEDCT batch jobs. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum individual file size for batch jobs | Each supported Region: 40 Kilobytes | No | The maximum individual file size for batch jobs. |
| Minimum size of batch jobs (all files) | Each supported Region: 1 Bytes | No | The minimum size of batch jobs (all files). |
| Ontology (ICD10, RxNNorm, SNOMEDCT) Linking maximum batch job size (sum total for all the files submitted in a batch job). | Each supported Region: 1 Gigabytes | No | Ontology (ICD10, RxNNorm, SNOMEDCT) Linking maximum batch job size (sum total for all the files submitted in a batch job). |
| Transactions per second (TPS) for the DescribeEntitiesDetectionV2Job operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the DescribeEntitiesDetectionV2Job operation. |
| Transactions per second (TPS) for the DescribeICD10CMInferenceJob operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the DescribeICD10CMInferenceJob operation. |
| Transactions per second (TPS) for the DescribePHIDetectionJob operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the DescribePHIDetectionJob operation. |
| Transactions per second (TPS) for the DescribeRxNormInferenceJob operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the DescribeRxNormInferenceJob operation. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the DetectEntities operation | Each supported Region: 40 | No | The maximum transactions per second (TPS) for the DetectEntities operation. |
| Transactions per second (TPS) for the DetectEntities-v2 operation | Each supported Region: 40 | No | The maximum transactions per second (TPS) for the DetectEntities-v2 operation. |
| Transactions per second (TPS) for the DetectPHI operation | Each supported Region: 40 | No | The maximum transactions per second (TPS) for the DetectPHI operation. |
| Transactions per second (TPS) for the InferICD10CM operation | Each supported Region: 40 | No | The maximum transactions per second (TPS) for the InferICD10CM operation. |
| Transactions per second (TPS) for the InferRxNorm operation | Each supported Region: 40 | No | The maximum transactions per second (TPS) for the InferRxNorm operation. |
| Transactions per second (TPS) for the InferSNOMEDCT operation | Each supported Region: 2 | No | The maximum transactions per second (TPS) for the InferSNOMEDCT operation. |
| Transactions per second (TPS) for the ListEntitiesDetectionV2Jobs operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the ListEntitiesDetectionV2Jobs operation. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the ListICD10CMInferenceJobs operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the ListICD10CMInferenceJobs operation. |
| Transactions per second (TPS) for the ListPHIDetectionJobs operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the ListPHIDetectionJobs operation. |
| Transactions per second (TPS) for the ListRxNormInferenceJobs operation | Each supported Region: 10 | Yes | The maximum transactions per second (TPS) for the ListRxNormInferenceJobs operation. |
| Transactions per second (TPS) for the StartEntitiesDetectionV2Job operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StartEntitiesDetectionV2Job operation. |
| Transactions per second (TPS) for the StartICD10CMInferenceJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StartICD10CMInferenceJob operation. |
| Transactions per second (TPS) for the StartPHIDetectionJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StartPHIDetectionJob operation. |
| Transactions per second (TPS) for the StartRxNormInferenceJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StartRxNormInferenceJob operation. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second (TPS) for the StartSNOMEDCTInferenceJob operation | Each supported Region: 5 | No | The maximum transactions per second (TPS) for the StartSNOMEDCTInferenceJob operation. |
| Transactions per second (TPS) for the StopEntitiesDetectionV2Job operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StopEntitiesDetectionV2Job operation. |
| Transactions per second (TPS) for the StopICD10CMInferenceJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StopICD10CMInferenceJob operation. |
| Transactions per second (TPS) for the StopPHIDetectionJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StopPHIDetectionJob operation. |
| Transactions per second (TPS) for the StopRxNormInferenceJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StopRxNormInferenceJob operation. |
| Transactions per second (TPS) for the StopSNOMEDCTInferenceJob operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StopSNOMEDCTInferenceJob operation. |

# AWS Compute Optimizer endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | compute-optimizer.us-east-2.amazonaws.com<br>compute-optimizer.us-east-2.api.aws | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | compute-optimizer.us-east-1.amazonaws.com<br>compute-optimizer.us-east-1.api.aws | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | compute-optimizer.us-west-1.amazonaws.com<br>compute-optimizer.us-west-1.api.aws | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | compute-optimizer.us-west-2.amazonaws.com<br>compute-optimizer.us-west-2.api.aws | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | compute-optimizer.af-south-1.amazonaws.com<br>compute-optimizer.af-south-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | compute-optimizer.ap-east-1.amazonaws.com<br>compute-optimizer.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | compute-optimizer.ap-south-2.amazonaws.com<br>compute-optimizer.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | compute-optimizer.ap-southeast-3.amazonaws.com<br>compute-optimizer.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | compute-optimizer.ap-southeast-4.amazonaws.com<br>compute-optimizer.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | compute-optimizer.ap-south-1.amazonaws.com<br>compute-optimizer.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | compute-optimizer.ap-northeast-3.amazonaws.com<br>compute-optimizer.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | compute-optimizer.ap-northeast-2.amazonaws.com<br>compute-optimizer.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | compute-optimizer.ap-southeast-1.amazonaws.com<br>compute-optimizer.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | compute-optimizer.ap-southeast-2.amazonaws.com<br>compute-optimizer.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | compute-optimizer.ap-northeast-1.amazonaws.com<br>compute-optimizer.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | compute-optimizer.ca-central-1.amazonaws.com<br>compute-optimizer.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | compute-optimizer.eu-central-1.amazonaws.com<br>compute-optimizer.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | compute-optimizer.eu-west-1.amazonaws.com<br>compute-optimizer.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | compute-optimizer.eu-west-2.amazonaws.com<br>compute-optimizer.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | compute-optimizer.eu-south-1.amazonaws.com<br>compute-optimizer.eu-south-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | compute-optimizer.eu-west-3.amazonaws.com<br>compute-optimizer.eu-west-3.api.aws | HTTPS<br>HTTPS | |
| Europe (Spain) | eu-south-2 | compute-optimizer.eu-south-2.amazonaws.com<br>compute-optimizer.eu-south-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | compute-optimizer.eu-north-1.amazonaws.com<br>compute-optimizer.eu-north-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Zurich) | eu-central-2 | compute-optimizer.eu-central-2.amazonaws.com<br>compute-optimizer.eu-central-2.api.aws | HTTPS<br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | compute-optimizer.il-central-1.amazonaws.com<br>compute-optimizer.il-central-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | compute-optimizer.me-south-1.amazonaws.com<br>compute-optimizer.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | compute-optimizer.me-central-1.amazonaws.com<br>compute-optimizer.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | compute-optimizer.sa-east-1.amazonaws.com<br>compute-optimizer.sa-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | compute-optimizer-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | compute-optimizer-fips.us-gov-west-1.amazonaws.com | HTTPS | |

**Note**

AWS Compute Optimizer offers dual stack endpoints so that you can access the service using IPv4 and IPv6 requests. For more information, see Dual stack endpoints.

### Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| The number of API calls per second per account | Each supported Region: 5 | No | The number of API calls per second per account. |

## AWS Config endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | config.us-east-2.amazonaws.com<br>config-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | config.us-east-1.amazonaws.com<br>config-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | config.us-west-1.amazonaws.com<br>config-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | config.us-west-2.amazonaws.com<br>config-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | config.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | config.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | config.ap-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | config.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | config.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | config.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | config.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | config.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | config.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | config.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | config.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | config.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | config.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | config.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | config.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | config.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | config.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | config.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | config.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | config.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | config.eu-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | config.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | config.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | config.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | config.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | config.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | config.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | config.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | config.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | config.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | config.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | config.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

### Resource tags

| Name | Default | Adjustable |
|---|---|---|
| Maximum number of tags per resource | 50 | No |

### AWS Config rules

| Name | Default | Adjustable |
|---|---|---|
| Maximum number of AWS Config Rules per Region per account | 1000 | No |

### Single Account Conformance Packs

| Name | Default | Adjustable |
|---|---|---|
| Maximum number of conformance packs per account | 50 | No |
| Maximum number of AWS Config Rules per conformance pack | 130 | No |

**Note**

AWS Config rules in conformance packs count in the quota for the Maximum number of AWS Config Rules per Region per account.

### Organization Conformance Packs

| Name | Default | Adjustable |
|---|---|---|
| Maximum number of conformance packs per organization | 50 | No |
| Maximum number of AWS Config Rules per organization conformance pack | 130 | No |

**Note**

Deploying at the organization level counts in quota for child accounts. AWS Config rules in conformance packs count in the quota for the Maximum number of AWS Config Rules per Region per account.

### Aggregators

| Name | Default | Adjustable |
|---|---|---|
| Maximum number of configuration aggregators | 50 | Yes |
| Maximum number of accounts in an aggregator | 10000 | No |
| Maximum number of accounts added or deleted per week for all aggregators | 1000 | Yes |

> **Note**
>
> Organization level aggregators and individual account aggregators both count in the quota for the Maximum number of configuration aggregators.

**Advanced queries**

| Name | Default | Adjustable |
| --- | --- | --- |
| Maximum number of saved queries in a single account and a Region | 300 | Yes |

# Amazon Connect Customer endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | connect.us-east-1.amazonaws.com | HTTPS |
| | | connect-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | connect.us-west-2.amazonaws.com | HTTPS |
| | | connect-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | connect.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | connect.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | connect.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | connect.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | connect.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | connect.ca-central-1.amazonaws.com<br>connect-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | connect.eu-central-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | connect.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | connect.us-gov-west-1.amazonaws.com | HTTPS | |

## Connect Customer Contact Lens endpoints

The Connect Customer Contact Lens service has the following endpoints.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | contact-lens.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | contact-lens.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | contact-lens.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | contact-lens.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | contact-lens.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | contact-lens.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | contact-lens.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | contact-lens.ca-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | contact-lens.eu-central-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | contact-lens.eu-west-2.amazonaws.com | HTTPS | |

## Connect Customer Participant Service endpoints

The Connect Customer Participant Service has the following endpoints.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | participant.connect.us-east-1.amazonaws.com | HTTPS |
| | | participant.connect-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | participant.connect.us-west-2.amazonaws.com | HTTPS |
| | | participant.connect-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | participant.connect.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | participant.connect.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | participant.connect.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | participant.connect.ap-southeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | participant.connect.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | participant.connect.ca-central-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | participant.connect.eu-central-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | participant.connect.eu-west-2.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | participant.connect.us-gov-west-1.amazonaws.com | HTTPS |

## Connect Customer Customer Profiles endpoints

The Connect Customer Customer Profiles service has the following endpoints.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | profile.us-east-1.amazonaws.com | HTTPS | |
| | | profile-fips.us-east-1.api.aws | HTTPS | |
| | | profile-fips.us-east-1.amazonaws.com | HTTPS | |
| | | profile.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | profile.us-west-2.amazonaws.com | HTTPS | |
| | | profile-fips.us-west-2.api.aws | HTTPS | |
| | | profile.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | profile-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | profile.af-south-1.amazonaws.com | HTTPS |
| | | profile.af-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | profile.ap-northeast-2.amazonaws.com | HTTPS |
| | | profile.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | profile.ap-southeast-1.amazonaws.com | HTTPS |
| | | profile.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | profile.ap-southeast-2.amazonaws.com | HTTPS |
| | | profile.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | profile.ap-northeast-1.amazonaws.com | HTTPS |
| | | profile.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | profile.ca-central-1.amazonaws.com | HTTPS |
| | | profile.ca-central-1.api.aws | HTTPS |
| | | profile-fips.ca-central-1.amazonaws.com | HTTPS |
| | | profile-fips.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | profile.eu-central-1.amazonaws.com | HTTPS |
| | | profile.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | profile.eu-west-2.amazonaws.com<br><br>profile.eu-west-2.api.aws | HTTPS<br><br>HTTPS | |

## AppIntegrations Service endpoints

The AppIntegrations Service has the following endpoints.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | app-integrations.us-east-1.amazonaws.com<br><br>app-integrations-fips.us-east-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| US West (Oregon) | us-west-2 | app-integrations.us-west-2.amazonaws.com<br><br>app-integrations-fips.us-west-2.amazonaws.com | HTTPS<br><br>HTTPS | |
| Africa (Cape Town) | af-south-1 | app-integrations.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | app-integrations.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | app-integrations.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | app-integrations.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | app-integrations.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | app-integrations.ca-central-1.amazonaws.com<br>app-integrations-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | app-integrations.eu-central-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | app-integrations.eu-west-2.amazonaws.com | HTTPS |

## Connect Customer Cases endpoints

Connect Customer Cases has the following endpoints.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | cases.us-east-1.amazonaws.com<br>cases-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | cases.us-west-2.amazonaws.com<br>cases-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | cases.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cases.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | cases.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cases.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cases.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | cases.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | cases.eu-central-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | cases.eu-west-2.amazonaws.com | HTTPS | |

## Connect Customer Voice ID endpoints

Connect Customer Voice ID has the following endpoints.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | voiceid.us-east-1.amazonaws.com<br>voiceid-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | voiceid.us-west-2.amazonaws.com<br>voiceid-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | voiceid.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | voiceid.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | voiceid.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | voiceid.ca-central-1.amazonaws.com<br>voiceid-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | voiceid.eu-central-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | voiceid.eu-west-2.amazonaws.com | HTTPS |

## Amazon Q in Connect endpoints

Amazon Q in Connect has the following endpoints.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | wisdom.us-east-1.amazonaws.com<br>ui.wisdom.us-east-1.prod.internal.clover.aws.dev<br>wisdom-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | wisdom.us-west-2.amazonaws.com<br>ui.wisdom.us-west-2.prod.internal.clover.aws.dev<br>wisdom-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | wisdom.af-south-1.amazonaws.com<br>ui.wisdom.af-south-1.prod.internal.clover.aws.dev | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | wisdom.ap-northeast-2.amazonaws.com<br>ui.wisdom.ap-northeast-2.prod.internal.clover.aws.dev | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | wisdom.ap-southeast-1.amazonaws.com<br>ui.wisdom.ap-southeast-1.prod.internal.clover.aws.dev | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | wisdom.ap-southeast-2.amazonaws.com<br>ui.wisdom.ap-southeast-2.prod.internal.clover.aws.dev | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | wisdom.ap-northeast-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | ui.wisdom.ap-northeast-1.prod.internal.clover.aws.dev | | |
| Canada (Central) | ca-central-1 | wisdom.ca-central-1.amazonaws.com<br>ui.wisdom.ca-central-1.prod.internal.clover.aws.dev<br>wisdom-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | wisdom.eu-central-1.amazonaws.com<br>ui.wisdom.eu-central-1.prod.internal.clover.aws.dev | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | wisdom.eu-west-2.amazonaws.com<br>ui.wisdom.eu-west-2.prod.internal.clover.aws.dev | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AWS Lambda functions per instance | Each supported Region: 50 | Yes | The maximum number of AWS Lambda functions you can create in this instance in the current Region. |
| Additional email addresses per queue | Each supported Region: 49 | Yes | The maximum number of additional email addresses that can be associated with a queue |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | in this instance in the current Region. |
| Agent status per instance | Each supported Region: 50 | No | The maximum number of agent statuses you can create in this instance in the current Region. This limit cannot be increased . |
| Agents per schedule | Each supported Region: 5,000 | No | The maximum number of agents per schedule in the current region. |
| Agents per staffing group | Each supported Region: 350 | No | The maximum number of agents per staffing group in the current region. |
| Amazon Connect instance count | Each supported Region: 2 | Yes | The maximum number of Amazon Connect instances you can create in this account in the current Region. |
| Amazon Lex V2 bot aliases per instance | Each supported Region: 100 | Yes | The maximum number of Amazon Lex V2 bot aliases you can use in this instance |
| Amazon Lex bots per instance | Each supported Region: 70 | Yes | The maximum number of Amazon Lex bots you can use in this instance in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amazon Pinpoint application integration associations per instance | Each supported Region: 1 | No | The maximum number of Amazon Pinpoint application integration associations you can create in this instance in the current Region. |
| Application integration associations per instance | Each supported Region: 25 | Yes | The maximum number of application integration associations you can create in this instance in the current Region. |
| Attributes per data table | Each supported Region: 100 | No | The maximum number of attributes that can be defined per data table. |
| Capacity plan override uploads per instance | Each supported Region: 5,000 | No | The maximum number of capacity plan override uploads per instance in the current region. |
| Capacity plan user data uploads per instance | Each supported Region: 500 | No | The maximum number of capacity plan user data uploads per instance in the current region. |
| Capacity planning scenarios per instance | Each supported Region: 500 | No | The maximum number of capacity planning scenarios you can create in this instance in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Capacity plans per instance | Each supported Region: 500 | No | The maximum number of capacity plans you can create in this instance in the current region. |
| Cases domain integration associations per instance | Each supported Region: 1 | No | The maximum number of Cases domain integration associations you can create in this instance in the current Region. |
| Concurrent active calls per instance | Each supported Region: 10 | Yes | The maximum number of concurrent active calls you can have in this instance in the current Region. If this is exceeded, contacts will get a fast busy tone, which indicates the transmission path to the called number is not available. |
| Concurrent active chats per instance | Each supported Region: 500 | Yes | The maximum number of concurrent active chats you can have in this instance in the current Region. If this is exceeded, additional chat sessions cannot be initiated. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent active emails per instance | Each supported Region: 1,000 | Yes | The maximum number of concurrent active emails you can have in this instance in the current Region. If this is exceeded, additional emails cannot be created. |
| Concurrent active tasks per instance | Each supported Region: 2,500 | Yes | The maximum number of concurrent active tasks you can have in this instance in the current Region. If this is exceeded, additional tasks cannot be created. |
| Concurrent campaign active calls per instance | Each supported Region: 0 | Yes | The maximum number of concurrent campaign active calls you can have in this instance in the current Region. If this is exceeded, contacts will get a fast busy tone, which indicates the transmission path to the called number is not available. |
| Concurrent uploads per instance | Each supported Region: 20 | No | The maximum number of uploads able to be processed at a time for an instance in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Connect AI agent assistant integration associations per instance | Each supported Region: 1 | No | The maximum number of Connect AI agent assistant integration associations you can create in this instance in the current Region. |
| Connect AI agent knowledge base integration associations per instance | Each supported Region: 10 | No | The maximum number of Connect AI agent knowledge base integration associations you can create in this instance in the current Region. |
| Connect AI agent message templates integration associations per instance | Each supported Region: 1 | No | The maximum number of Connect AI agent message templates integration associations you can create in this instance in the current Region. |
| Connect AI agent quick responses integration associations per instance | Each supported Region: 1 | No | The maximum number of Connect AI agent quick responses integration associations you can create in this instance in the current Region. |
| Contact Lens connectors per account | Each supported Region: 0 | Yes | The maximum number of Contact Lens connectors that can be created in this AWS account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Contact flows per instance | Each supported Region: 100 | Yes | The maximum number of contact flows you can create in this instance in the current Region. |
| Data tables per instance | Each supported Region: 100 | Yes | The maximum number of data tables that can be created per Amazon Connect instance. |
| Demand groups per forecast group | Each supported Region: 25 | No | The maximum number of demand groups you can have in a given forecast group. |
| Email addresses per instance | Each supported Region: 100 | Yes | The maximum number of email addresses you can have in this instance in the current Region. |
| Event integration associations per instance | Each supported Region: 10 | Yes | The maximum number of event integration associations you can create in this instance in the current Region. |
| External voice transfer connectors per account | Each supported Region: 0 | Yes | The maximum number of external voice transfer connectors that can be created in this AWS account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| File scanner integration associations per instance | Each supported Region: 1 | No | The maximum number of file scanner integration associations you can create in this instance in the current Region. |
| File size per upload of capacity plan overrides | Each supported Region: 250 | No | The maximum number of megabytes a capacity plan override file upload can be for an instance in the current region. |
| File size per upload of capacity plan user data | Each supported Region: 1,000 | No | The maximum number of megabytes a capacity plan user data file upload can be for an instance in the current region. |
| File size per upload of forecast overrides | Each supported Region: 250 | No | The maximum number of megabytes a forecast override file upload can be for an instance in the current region. |
| File size per upload of historical actuals | Each supported Region: 1,000 | No | The maximum number of megabytes a historical actuals upload can be for an instance in the current region. |
| Flow modules per instance | Each supported Region: 200 | Yes | The maximum number of flow modules you can create per instance in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Forecast groups per instance | Each supported Region: 500 | No | The maximum number of forecast groups you can create in this instance in the current region. |
| Forecast override uploads per instance | Each supported Region: 500 | No | The maximum number of forecast override uploads per instance in the current region. |
| Historical actuals 15 or 30 minute interval aggregated file size limit | Each supported Region: 2,000 | No | The maximum total size in megabytes for all historical actuals files with 15 or 30 minute interval aggregated together for an instance in the current region. |
| Historical actuals 15 or 30 minute interval file count | Each supported Region: 300 | No | The maximum number of historical actuals files with 15 or 30 minute interval per instance in the current region. |
| Historical actuals daily interval aggregated file size limit | Each supported Region: 2,000 | No | The maximum total size in megabytes for all historical actuals files with daily interval aggregated together for an instance in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Historical actuals daily interval file count | Each supported Region: 300 | No | The maximum number of historical actuals files with daily interval per instance in the current region. |
| Hours of operation overrides per instance | Each supported Region: 50 | No | The maximum number of overrides per hours of operation in this instance in the current Region. This limit cannot be increased. |
| Hours of operation per instance | Each supported Region: 100 | Yes | The maximum number of hours of operation you can create in this instance in the current Region. |
| Maximum active recording sessions from external voice systems per instance | Each supported Region: 10 | Yes | The maximum number of active recording sessions from external voice systems you can have per Amazon Connect instance in the current Region. If this is exceeded, additional analytics sessions cannot be started. |
| Notifications per instance | Each supported Region: 100 | Yes | The maximum number of notifications you can create in this instance in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Phone numbers per instance | Each supported Region: 5 | Yes | The maximum number of phone numbers you can claim for this instance in the current Region. |
| Predefined Attributes | Each supported Region: 150 | Yes | The maximum number of predefined attributes you can create per instance in the current Region. |
| Primary attributes per data table | Each supported Region: 5 | No | The maximum number of primary (key) attributes that can be defined per data table. |
| Proficiencies per agent | Each supported Region: 20 | Yes | The maximum number of proficiencies you can associate with an agent in the current Region. |
| Prompts per instance | Each supported Region: 500 | Yes | The maximum number of prompts you can create in this instance in the current Region. |
| Queues per forecast group | Each supported Region: 200 | No | The maximum number of queues you can have in a given forecast group. |
| Queues per instance | Each supported Region: 100 | Yes | The maximum number of queues you can create in this instance in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Queues per routing profile per instance | Each supported Region: 50 | Yes | The maximum number of queues you can create per routing profile in this instance in the current Region. |
| Quick connects per instance | Each supported Region: 100 | Yes | The maximum number of quick connects/transfer destinations you can create in this instance in the current Region. |
| Rate of AssociateApprovedOrigin API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateApprovedOrigin API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateBot API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateBot API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of AssociateContactWithUser API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateContactWithUser API requests allowed per second. When you reach this quota, Amazon Connect rejects additional requests for this operation for the remainder of the interval. |
| Rate of AssociateDefaultVocabulary API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateDefaultVocabulary API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateHoursOfOperations API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateHoursOfOperations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of AssociateInstanceStorageConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateInstanceStorageConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateLambdaFunction API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateLambdaFunction API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateLexBot API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateLexBot API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of AssociatePhoneNumberContactFlow API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociatePhoneNumberContactFlow API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateQueueEmailAddresses API requests | Each supported Region: 5 per second | Yes | The maximum number of AssociateQueueEmailAddresses API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateQueueQuickConnects API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateQueueQuickConnects API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of AssociateRoutingProfileQueues API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateRoutingProfileQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateSecurityKey API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateSecurityKey API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateSecurityProfiles API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateSecurityProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects additional requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of AssociateUserProficiencies API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateUserProficiencies API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of AssociateWorkspace API requests | Each supported Region: 2 per second | Yes | The maximum number of AssociateWorkspace API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of BatchCreateDataTableValue API requests | Each supported Region: 2 per second | Yes | The maximum number of BatchCreateDataTableValue API requests allowed per second. |
| Rate of BatchDeleteDataTableValue API requests | Each supported Region: 2 per second | Yes | The maximum number of BatchDeleteDataTableValue API requests allowed per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of BatchDescribeDataTableValue API requests | Each supported Region: 2 per second | Yes | The maximum number of BatchDescribeDataTableValue API requests allowed per second. |
| Rate of BatchPutContact API requests | Each supported Region: 10 per second | Yes | The maximum number of BatchPutContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of BatchUpdateDataTableValue API requests | Each supported Region: 2 per second | Yes | The maximum number of BatchUpdateDataTableValue API requests allowed per second. |
| Rate of ClaimPhoneNumber API requests | Each supported Region: 2 per second | Yes | The maximum number of ClaimPhoneNumber API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateAgentStatus API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateAgentStatus API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateAttachedFile API requests | Each supported Region: 0 per second | Yes | The maximum number of CreateAttachedFile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateContactFlow API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateContactFlow API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateContactFlowModule API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateContactFlowModule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateDataTable API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateDataTable API requests allowed per second. |
| Rate of CreateDataTableAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateDataTableAttribute API requests allowed per second. |
| Rate of CreateEmailAddress API requests | Each supported Region: 1 per second | Yes | The maximum number of CreateEmailAddress API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateHoursOfOperation API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateHoursOfOperation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateHoursOfOperationOverride API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateHoursOfOperationOverride API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateInstance API requests | Each supported Region: 0.1 per second | Yes | The maximum number of CreateInstance API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateIntegrationAssociation API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateIntegrationAssociation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateNotification API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateNotification API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateParticipant API requests | Each supported Region: 5 per second | Yes | The maximum number of CreateParticipant API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreatePredefinedAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of CreatePredefinedAttribute API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreatePrompt API requests | Each supported Region: 2 per second | Yes | The maximum number of CreatePrompt API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreatePushNotificationRegistration API requests | Each supported Region: 5 per second | Yes | The maximum number of CreatePushNotificationRegistration API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateQueue API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateQueue API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateQuickConnect API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateQuickConnect API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateRoutingProfile API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateRoutingProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateRule API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateRule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateSecurityProfile API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateSecurityProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateTaskTemplate API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateTaskTemplate API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateTrafficDistributionGroup API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateTrafficDistributionGroup API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateUseCase API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateUseCase API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateUser API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateUser API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateUserHierarchyGroup API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateUserHierarchyGroup API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateVocabulary API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateVocabulary API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of CreateWorkspace API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateWorkspace API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateWorkspacePage API requests | Each supported Region: 2 per second | Yes | The maximum number of CreateWorkspacePage API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteContactFlow API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteContactFlow API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteContactFlowModule API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteContactFlowModule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteDataTable API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteDataTable API requests allowed per second. |
| Rate of DeleteDataTableAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteDataTableAttribute API requests allowed per second. |
| Rate of DeleteEmailAddress API requests | Each supported Region: 1 per second | Yes | The maximum number of DeleteEmailAddress API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteHoursOfOperation API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteHoursOfOperation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteHoursOfOperationOverride API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteHoursOfOperationOverride API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteInstance API requests | Each supported Region: 0.1 per second | Yes | The maximum number of DeleteInstance API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteIntegrationAssociation API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteIntegrationAssociation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteNotification API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteNotification API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeletePredefinedAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of DeletePredefinedAttribute API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeletePrompt API requests | Each supported Region: 2 per second | Yes | The maximum number of DeletePrompt API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeletePushNotificationRegistration API requests | Each supported Region: 5 per second | Yes | The maximum number of DeletePushNotificationRegistration API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteQueue API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteQueue API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteQuickConnect API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteQuickConnect API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteRule API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteRule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteSecurityProfile API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteSecurityProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteTaskTemplate API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteTaskTemplate API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteTrafficDistributionGroup API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteTrafficDistributionGroup API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteUseCase API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteUseCase API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteUser API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteUser API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteUserHierarchyGroup API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteUserHierarchyGroup API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteVocabulary API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteVocabulary API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteWorkspace API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteWorkspace API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteWorkspaceMedia API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteWorkspaceMedia API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DeleteWorkspacePage API requests | Each supported Region: 2 per second | Yes | The maximum number of DeleteWorkspacePage API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeAgentStatus API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeAgentStatus API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeContact API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeContactFlow API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeContactFlow API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeContactFlowModule API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeContactFlowModule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeDataTable API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeDataTable API requests allowed per second. |
| Rate of DescribeDataTableAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeDataTableAttribute API requests allowed per second. |
| Rate of DescribeEmailAddress API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeEmailAddress API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeHoursOfOperation API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeHoursOfOperation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeHoursOfOperationOverride API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeHoursOfOperationOverride API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeInstance API requests | Each supported Region: 1 per second | Yes | The maximum number of DescribeInstance API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeInstanceAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeInstanceAttribute API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeInstanceStorageConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeInstanceStorageConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeNotification API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeNotification API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribePhoneNumber API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribePhoneNumber API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribePredefinedAttribute API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribePredefined Attribute API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribePrompt API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribePrompt API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeQueue API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeQueue API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeQuickConnect API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeQuickConnect API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeRoutingProfile API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeRoutingProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeRule API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeRule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeSecurityProfile API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeSecurityProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeTrafficDistribution Group API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeTrafficDistributionGroup API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeUser API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeUser API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeUserHierarchyGroup API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeUserHierarchyGroup API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeUserHierarchyStructure API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeUserHierarchyStructure API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DescribeVocabulary API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeVocabulary API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeWorkspace API requests | Each supported Region: 2 per second | Yes | The maximum number of DescribeWorkspace API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateApprovedOrigin API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateApprovedOrigin API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateBot API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateBot API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DisassociateHoursOfOperations API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateHoursOfOperations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateInstanceStorageConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateInstanceStorageConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateLambdaFunction API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateLambdaFunction API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DisassociateLexBot API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateLexBot API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociatePhoneNumberContactFlow API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociatePhoneNumberContactFlow API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateQueueEmailAddresses API requests | Each supported Region: 5 per second | Yes | The maximum number of DisassociateQueueEmailAddresses API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DisassociateQueueQuickConnects API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateQueueQuickConnects API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateRoutingProfileQueues API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateRoutingProfileQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateSecurityKey API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateSecurityKey API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DisassociateSecurityProfiles API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateSecurityProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects additional requests for this operation for the remainder of the interval. |
| Rate of DisassociateUserProficiencies API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateUserProficiencies API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of DisassociateWorkspace API requests | Each supported Region: 2 per second | Yes | The maximum number of DisassociateWorkspace API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DismissUserContact API requests | Each supported Region: 2 per second | Yes | The maximum number of DismissUserContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of EvaluateDataTableValues API requests | Each supported Region: 2 per second | Yes | The maximum number of EvaluateDataTableValues API requests allowed per second. |
| Rate of GetAnalyticsInsights API requests | Each supported Region: 50 per second | Yes | The maximum number of GetAnalyticsInsights API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetContactAttributes API requests | Each supported Region: 10 per second | Yes | The maximum number of GetContactAttributes API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetContactMetrics API requests | Each supported Region: 2 per second | Yes | The maximum number of GetContactMetrics API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetCurrentMetricData API requests | Each supported Region: 5 per second | Yes | The maximum number of GetCurrentMetricData API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetCurrentUserData API requests | Each supported Region: 2 per second | Yes | The maximum number of GetCurrentUserData API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetEffectiveHoursOfOperations API requests | Each supported Region: 2 per second | Yes | The maximum number of GetEffectiveHoursOfOperations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetFederationToken API requests | Each supported Region: 2 per second | Yes | The maximum number of GetFederationToken API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetMetricData API requests | Each supported Region: 5 per second | Yes | The maximum number of GetMetricData API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetMetricDataV2 API requests | Each supported Region: 10 per second | Yes | The maximum number of GetMetricDataV2 API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetTaskTemplate API requests | Each supported Region: 2 per second | Yes | The maximum number of GetTaskTemplate API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of GetTrafficDistribution API requests | Each supported Region: 2 per second | Yes | The maximum number of GetTrafficDistribution API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ImportWorkspaceMedia API requests | Each supported Region: 2 per second | Yes | The maximum number of ImportWorkspaceMedia API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListAgentStatuses API requests | Each supported Region: 10 per second | Yes | The maximum number of ListAgentStatuses API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListApprovedOrigins API requests | Each supported Region: 2 per second | Yes | The maximum number of ListApprovedOrigins API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListBots API requests | Each supported Region: 2 per second | Yes | The maximum number of ListBots API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListChildHoursOfOperations API requests | Each supported Region: 10 per second | Yes | The maximum number of ListChildHoursOfOperations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListContactFlowModules API requests | Each supported Region: 2 per second | Yes | The maximum number of ListContactFlowModules API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListContactFlows API requests | Each supported Region: 2 per second | Yes | The maximum number of ListContactFlows API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListContactReferences API requests | Each supported Region: 10 per second | Yes | The maximum number of ListContactReferences API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListDataTableAttributes API requests | Each supported Region: 2 per second | Yes | The maximum number of ListDataTableAttributes API requests allowed per second. |
| Rate of ListDataTablePrimaryValues API requests | Each supported Region: 2 per second | Yes | The maximum number of ListDataTablePrimaryValues API requests allowed per second. |
| Rate of ListDataTableValues API requests | Each supported Region: 2 per second | Yes | The maximum number of ListDataTableValues API requests allowed per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListDataTables API requests | Each supported Region: 2 per second | Yes | The maximum number of ListDataTables API requests allowed per second. |
| Rate of ListDefaultVocabularies API requests | Each supported Region: 2 per second | Yes | The maximum number of ListDefaultVocabularies API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListEntitySecurityProfiles API requests | Each supported Region: 10 per second | Yes | The maximum number of ListEntitySecurityProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects additional requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListHoursOfOperationOverrides API requests | Each supported Region: 2 per second | Yes | The maximum number of ListHoursOfOperationOverrides API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListHoursOfOperations API requests | Each supported Region: 10 per second | Yes | The maximum number of ListHoursOfOperations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListInstanceAttributes API requests | Each supported Region: 2 per second | Yes | The maximum number of ListInstanceAttributes API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListInstanceStorageConfigs API requests | Each supported Region: 2 per second | Yes | The maximum number of ListInstanceStorageConfigs API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListInstances API requests | Each supported Region: 1 per second | Yes | The maximum number of ListInstances API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListIntegrationAssociations API requests | Each supported Region: 50 per second | Yes | The maximum number of ListIntegrationAssociations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListLambdaFunctions API requests | Each supported Region: 2 per second | Yes | The maximum number of ListLambdaFunctions API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListLexBots API requests | Each supported Region: 2 per second | Yes | The maximum number of ListLexBots API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListNotifications API requests | Each supported Region: 2 per second | Yes | The maximum number of ListNotifications API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListPhoneNumbers API requests | Each supported Region: 2 per second | Yes | The maximum number of ListPhoneNumbers API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListPhoneNumbersV2 API requests | Each supported Region: 2 per second | Yes | The maximum number of ListPhoneNumbersV2 API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListPredefinedAttributes API requests | Each supported Region: 10 per second | Yes | The maximum number of ListPredefinedAttributes API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListPrompts API requests | Each supported Region: 10 per second | Yes | The maximum number of ListPrompts API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListQueueEmailAddresses API requests | Each supported Region: 15 per second | Yes | The maximum number of ListQueueEmailAddresses API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListQueueQuickConnects API requests | Each supported Region: 10 per second | Yes | The maximum number of ListQueueQuickConnects API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListQueues API requests | Each supported Region: 10 per second | Yes | The maximum number of ListQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListQuickConnects API requests | Each supported Region: 10 per second | Yes | The maximum number of ListQuickConnects API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListRoutingProfileManualAssignmentQueues API requests | Each supported Region: 2 per second | Yes | The maximum number of ListRoutingProfileManualAssignmentQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListRoutingProfileQueues API requests | Each supported Region: 10 per second | Yes | The maximum number of ListRoutingProfileQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListRoutingProfiles API requests | Each supported Region: 10 per second | Yes | The maximum number of ListRoutingProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListRules API requests | Each supported Region: 2 per second | Yes | The maximum number of ListRules API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListSecurityKeys API requests | Each supported Region: 2 per second | Yes | The maximum number of ListSecurityKeys API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListSecurityProfileFlowModules API requests | Each supported Region: 10 per second | Yes | The maximum number of ListSecurityProfileFlowModules API requests allowed per second. When you reach this quota, Amazon Connect rejects additional requests for this operation for the remainder of the interval. |
| Rate of ListSecurityProfilePermissions API requests | Each supported Region: 2 per second | Yes | The maximum number of ListSecurityProfilePermissions API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListSecurityProfiles API requests | Each supported Region: 2 per second | Yes | The maximum number of ListSecurityProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListTagsForResource API requests | Each supported Region: 2 per second | Yes | The maximum number of ListTagsForResource API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListTaskTemplates API requests | Each supported Region: 2 per second | Yes | The maximum number of ListTaskTemplates API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListTrafficDistributionGroups API requests | Each supported Region: 2 per second | Yes | The maximum number of ListTrafficDistributionGroups API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListUseCases API requests | Each supported Region: 2 per second | Yes | The maximum number of ListUseCases API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListUserHierarchyGroups API requests | Each supported Region: 10 per second | Yes | The maximum number of ListUserHierarchyGroups API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListUserNotifications API requests | Each supported Region: 2 per second | Yes | The maximum number of ListUserNotifications API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListUserProficiencies API requests | Each supported Region: 10 per second | Yes | The maximum number of ListUserProficiencies API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListUsers API requests | Each supported Region: 10 per second | Yes | The maximum number of ListUsers API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListWorkspaceMedia API requests | Each supported Region: 2 per second | Yes | The maximum number of ListWorkspaceMedia API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListWorkspacePages API requests | Each supported Region: 2 per second | Yes | The maximum number of ListWorkspacePages API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ListWorkspaces API requests | Each supported Region: 2 per second | Yes | The maximum number of ListWorkspaces API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of MonitorContact API requests | Each supported Region: 2 per second | Yes | The maximum number of MonitorContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of PauseContact API requests | Each supported Region: 2 per second | Yes | The maximum number of PauseContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of PutUserStatus API requests | Each supported Region: 2 per second | Yes | The maximum number of PutUserStatus API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ReleasePhoneNumber API requests | Each supported Region: 2 per second | Yes | The maximum number of ReleasePhoneNumber API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ReplicateInstance API requests | Each supported Region: 2 per second | Yes | The maximum number of ReplicateInstance API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of ResumeContact API requests | Each supported Region: 2 per second | Yes | The maximum number of ResumeContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ResumeContactRecording API requests | Each supported Region: 2 per second | Yes | The maximum number of ResumeContactRecording API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchAgentStatuses API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchAgentStatuses API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchAvailablePhoneNumbers API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchAvailablePhoneNumbers API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchDataTables API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchDataTables API requests allowed per second. |
| Rate of SearchEmailAddresses API requests | Each supported Region: 15 per second | Yes | The maximum number of SearchEmailAddresses API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchHoursOfOperationOverrides API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchHoursOfOperationOverrides API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchHoursOfOperations API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchHoursOfOperations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchNotifications API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchNotifications API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchPredefinedAttributes API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchPredefinedAttributes API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchPrompts API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchPrompts API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchQueues API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchQuickConnects API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchQuickConnects API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchResourceTags API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchResourceTags API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchRoutingProfiles API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchRoutingProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchRules API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchRules API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchSecurityProfiles API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchSecurityProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchUserHierarchyGroups API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchUserHierarchyGroups API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchUsers API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchUsers API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchVocabularies API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchVocabularies API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchWorkspaceAssociations API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchWorkspaceAssociations API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SearchWorkspaces API requests | Each supported Region: 2 per second | Yes | The maximum number of SearchWorkspaces API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SendIntegrationEvent API requests | Each supported Region: 10 per second | Yes | The maximum number of SendIntegrationEvent API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SendOutboundChatMessage API requests | Each supported Region: 100 per second | Yes | The maximum number of SendOutboundChatMessage API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SendOutboundEmail API requests | Each supported Region: 2 per second | Yes | The maximum number of SendOutboundEmail API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartChatContact API requests | Each supported Region: 5 per second | Yes | The maximum number of StartChatContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartContactConversationalAnalyticsJob API requests | Each supported Region: 0 per second | Yes | The maximum number of StartContactConversationalAnalyticsJob API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartContactRecording API requests | Each supported Region: 2 per second | Yes | The maximum number of StartContactRecording API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartContactStreaming API requests | Each supported Region: 5 per second | Yes | The maximum number of StartContactStreaming API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartEmailContact API requests | Each supported Region: 2 per second | Yes | The maximum number of StartEmailContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartOutboundChatContact API requests | Each supported Region: 2 per second | Yes | The maximum number of StartOutboundChatContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartOutboundEmailContact API requests | Each supported Region: 2 per second | Yes | The maximum number of StartOutboundEmailContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartOutboundVoiceContact API requests | Each supported Region: 2 per second | Yes | The maximum number of StartOutboundVoiceContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartScreenSharing API requests | Each supported Region: 2 per second | Yes | The maximum number of StartScreenSharing API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartTaskContact API requests | Each supported Region: 2 per second | Yes | The maximum number of StartTaskContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StartWebRTCContact API requests | Each supported Region: 2 per second | Yes | The maximum number of StartWebRTCContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StopContact API requests | Each supported Region: 10 per second | Yes | The maximum number of StopContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StopContactRecording API requests | Each supported Region: 2 per second | Yes | The maximum number of StopContactRecording API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of StopContactStreaming API requests | Each supported Region: 5 per second | Yes | The maximum number of StopContactStreaming API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of SuspendContactRecording API requests | Each supported Region: 2 per second | Yes | The maximum number of SuspendContactRecording API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of TagContact API requests | Each supported Region: 20 per second | Yes | The maximum number of TagContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of TagResource API requests | Each supported Region: 2 per second | Yes | The maximum number of TagResource API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of TransferContact API requests | Each supported Region: 2 per second | Yes | The maximum number of TransferContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UntagContact API requests | Each supported Region: 20 per second | Yes | The maximum number of UntagContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UntagResource API requests | Each supported Region: 2 per second | Yes | The maximum number of UntagResource API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateAgentStatus API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateAgentStatus API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateContact API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateContact API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateContactAttributes API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateContactAttributes API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateContactFlowContent API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateContactFlowContent API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateContactFlowMetadata API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateContactFlowMetadata API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateContactFlowModuleContent API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateContactFlowModuleContent API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateContactFlowModuleMetadata API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateContactFlowModuleMetadata API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateContactFlowName API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateContactFlowName API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateContactSchedule API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateContactSchedule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateDataTableAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateDataTableAttribute API requests allowed per second. |
| Rate of UpdateDataTableMetadata API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateDataTableMetadata API requests allowed per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateDataTablePrimaryValues API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateDataTablePrimaryValues API requests allowed per second. |
| Rate of UpdateEmailAddressMetadata API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateEmailAddressMetadata API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateHoursOfOperation API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateHoursOfOperation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateHoursOfOperationOverride API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateHoursOfOperationOverride API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateInstanceAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateInstanceAttribute API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateInstanceStorageConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateInstanceStorageConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateNotificationContent API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateNotificationContent API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateParticipantRoleConfig API requests | Each supported Region: 5 per second | Yes | The maximum number of UpdateParticipantRoleConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdatePhoneNumber API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdatePhoneNumber API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdatePredefinedAttribute API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdatePredefinedAttribute API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdatePrompt API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdatePrompt API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateQueueHoursOfOperation API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueueHoursOfOperation API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateQueueMaxContacts API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueueMaxContacts API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateQueueName API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueueName API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateQueueOutboundCallerConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueueOutboundCallerConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateQueueOutboundEmailConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueueOutboundEmailConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateQueueStatus API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueueStatus API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateQuickConnectConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQuickConnectConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateQuickConnectName API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateQuickConnectName API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateRoutingProfileConcurrency API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateRoutingProfileConcurrency API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateRoutingProfileDefaultOutboundQueue API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateRoutingProfileDefaultOutboundQueue API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateRoutingProfileName API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateRoutingProfileName API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateRoutingProfileQueues API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateRoutingProfileQueues API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateRule API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateRule API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateSecurityProfile API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateSecurityProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateTaskTemplate API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateTaskTemplate API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateTrafficDistribution API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateTrafficDistribution API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateUserConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateUserHierarchy API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserHierarchy API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateUserHierarchyGroupName API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserHierarchyGroupName API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateUserHierarchyStructure API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserHierarchyStructure API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateUserIdentityInfo API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserIdentityInfo API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateUserNotificationStatus API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserNotificationStatus API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateUserPhoneConfig API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserPhoneConfig API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateUserProficiencies API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserProficiencies API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateUserRoutingProfile API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserRoutingProfile API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateUserSecurityProfiles API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateUserSecurityProfiles API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateWorkspaceMetadata API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateWorkspaceMetadata API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateWorkspacePage API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateWorkspacePage API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateWorkspaceTheme API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateWorkspaceTheme API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Rate of UpdateWorkspaceVisibility API requests | Each supported Region: 2 per second | Yes | The maximum number of UpdateWorkspaceVisibility API requests allowed per second. When you reach this quota, Amazon Connect rejects requests for this operation for the remainder of the interval. |
| Reports per instance | Each supported Region: 2,000 | Yes | The maximum number of reports that you can create in this instance in the current Region. Personal saved reports count towards the limit. As a best practice, we recommend you implement appropriate retention policies so reports dont pile up. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Routing profiles per instance | Each supported Region: 500 | Yes | The maximum number of routing profiles you can create in this instance in the current Region. |
| Scheduled reports per instance | Each supported Region: 100 | Yes | The maximum number of scheduled reports you can create in this instance in the current Region. |
| Schedules per instance | Each supported Region: 1,000 | No | The maximum number of schedules per instance in the current region. |
| Security profiles per instance | Each supported Region: 100 | Yes | The maximum number of security profiles you can create in this instance in the current Region. |
| Shift activities per instance | Each supported Region: 500 | No | The maximum number of shift activities per instance in the current region. |
| Shift activities per shift profile | Each supported Region: 10 | No | The maximum number of shift activities in a shift profile in the current region. |
| Shift profiles per instance | Each supported Region: 2,500 | No | The maximum number of shift profiles per instance in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Shift rotation pattern steps per shift profile | Each supported Region: 1,300 | No | The maximum number of shift rotation pattern steps associated per shift profile in the current Region. |
| Shift rotation patterns per instance | Each supported Region: 1,300 | No | The maximum number of shift rotation patterns per instance in the current Region. |
| Staffing groups per Forecast group | Each supported Region: 300 | No | The maximum number of staffing groups per forecast group in the current region. |
| Staffing groups per instance | Each supported Region: 1,300 | No | The maximum number of staffing groups per instance in the current region. |
| Staffing groups per supervisor | Each supported Region: 250 | No | The maximum number of staffing groups per supervisor in the current region. |
| Supervisors per staffing group | Each supported Region: 100 | No | The maximum number of supervisors per staffing group in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| User hierarchy groups per instance | Each supported Region: 500 | Yes | The maximum number of user hierarchy groups you can create in this instance in the current Region. |
| Users per instance | Each supported Region: 500 | Yes | The maximum number of users you can create in this instance in the current Region. |
| Values per data table | Each supported Region: 1,000 | Yes | The maximum number of values that can be stored per data table. |
| Voice ID integration associations per instance | Each supported Region: 1 | No | The maximum number of Voice ID integration associations you can create in this instance in the current Region. |
| Workspaces per instance | Each supported Region: 20 | Yes | The maximum number of workspaces you can create in this instance in the current Region. |

## Connect Customer Customer Profiles quotas

Connect Customer Customer Profiles has the following quotas.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amazon Connect Customer Profiles domain count | Each supported Region: 100 | Yes | The maximum number of Amazon Connect Customer Profiles domains you can create in this account in the current AWS Region. |
| Keys per object type | Each supported Region: 10 | Yes | The maximum number of keys that can be defined per object type in the current AWS Region. |
| Maximum expiration in days | Each supported Region: 1,098 | Yes | The maximum expiration, in days, that can be defined for an object or profile in the current AWS Region. |
| Maximum number of calculated attributes per domain | Each supported Region: 100 | No | The maximum number of calculated attributes per domain in the current AWS Region. |
| Maximum number of domain object types per domain | Each supported Region: 1 | No | The maximum number of domain object types you can define per domain in the current AWS Region. |
| Maximum number of event stream per domain | Each supported Region: 1 | No | The maximum number of event stream per domain in the current AWS Region. |
| Maximum number of event triggers per domain | Each supported Region: 50 | Yes | The maximum number of event triggers per |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | domain in the current AWS Region. |
| Maximum number of integrations | Each supported Region: 50 | Yes | The maximum number of integrations per domain in the current AWS Region. |
| Maximum number of profile history records per profile | Each supported Region: 10,000 | Yes | The maximum number of profile history records per profile in the current AWS Region. |
| Maximum number of recommender schemas per domain | Each supported Region: 5 | Yes | The maximum number of recommender schemas you can define per domain in the current AWS Region. |
| Maximum number of recommenders per domain | Each supported Region: 10 | Yes | The maximum number of recommenders you can define per domain in the current AWS Region. |
| Maximum number of recommeneder filters per domain | Each supported Region: 30 | No | The maximum number of recommender filters you can define per domain in the current AWS Region. |
| Maximum number of segment snapshots per day | Each supported Region: 200 | Yes | The maximum number of segmentation snapshots that can be created per domain per day in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum size of all objects for a profile | Each supported Region: 51,200 Kilobytes | Yes | The total size of a profile, including all of its related objects, in the current AWS Region. |
| Object and profile maximum size | Each supported Region: 250 Kilobytes | No | The maximum size of a single profile or profile object in the current AWS Region. |
| Object types per domain | Each supported Region: 100 | Yes | The maximum number of object types you can define per domain in the current AWS Region. |
| Objects per profile | Each supported Region: 1,000 | Yes | The maximum number of objects that can be attached to a single profile in the current AWS Region. |

## Connect Customer Cases quotas

Connect Customer Cases has the following quotas.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Attached SLAs per case | Each supported Region: 30 | Yes | The maximum number of SLAs that can be attached to a case |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Attached files per case | Each supported Region: 50 | Yes | The maximum number of files that can be attached to a case |
| BatchGetCaseRule burst quota | Each supported Region: 25 per second | Yes | The maximum number of concurrent BatchGetCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| BatchGetCaseRule rate quota | Each supported Region: 8 per second | Yes | The maximum number of BatchGetCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BatchGetField burst quota | Each supported Region: 25 per second | Yes | The maximum number of concurrent BatchGetField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| BatchGetField rate quota | Each supported Region: 8 per second | Yes | The maximum number of BatchGetField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| BatchPutFieldOptions burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent BatchPutFieldOptions API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BatchPutFieldOptions rate quota | Each supported Region: 2 per second | Yes | The maximum number of BatchPutFieldOptions API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| Case fields per layout | Each supported Region: 100 | Yes | The maximum number of case fields in a case layout |
| Case rules per domain | Each supported Region: 100 | Yes | The maximum number of case rules per domain |
| Child field values per field options case rule | Each supported Region: 500 | Yes | The maximum number of child field values per field options case rule |
| CreateCase burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent CreateCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateCase rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateCaseRule burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent CreateCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateCaseRule rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateDomain burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent CreateDomain API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateDomain rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateDomain API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateField burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent CreateField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateField rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateLayout burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent CreateLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateLayout rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateRelatedItem burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent CreateRelatedItem API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateRelatedItem rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateRelatedItem API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| CreateTemplate burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent CreateTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateTemplate rate quota | Each supported Region: 2 per second | Yes | The maximum number of CreateTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteCase burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent DeleteCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteCase rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteCaseRule burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent DeleteCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteCaseRule rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteDomain burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent DeleteDomain API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteDomain rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteDomain API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteField burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent DeleteField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteField rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteLayout burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent DeleteLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteLayout rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteRelatedItem burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent DeleteRelatedItem API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteRelatedItem rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteRelatedItem API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteTemplate burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent DeleteTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| DeleteTemplate rate quota | Each supported Region: 2 per second | Yes | The maximum number of DeleteTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Domains | Each supported Region: 5 | Yes | The maximum number of Cases domains per AWS account |
| Field options case rules per template | Each supported Region: 10 | Yes | The maximum number of field options case rules per template |
| Field options per field | Each supported Region: 500 | Yes | The maximum number of field options per single select field in the Cases domain |
| Fields per domain | Each supported Region: 500 | Yes | The maximum number of fields in a Cases domain |
| Fields per related item | Each supported Region: 5 | Yes | The maximum number of fields per related item |
| GetCase burst quota | Each supported Region: 15 per second | Yes | The maximum number of concurrent GetCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetCase rate quota | Each supported Region: 4 per second | Yes | The maximum number of GetCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetCaseAuditEvents burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent GetCaseAuditEvents API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetCaseAuditEvents rate quota | Each supported Region: 2 per second | Yes | The maximum number of GetCaseAuditEvents API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetCaseEventConfiguration burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent GetCaseEventConfiguration API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetCaseEventConfiguration rate quota | Each supported Region: 2 per second | Yes | The maximum number of GetCaseEventConfiguration API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetDomain burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent GetDomain API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetDomain rate quota | Each supported Region: 2 per second | Yes | The maximum number of GetDomain API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetLayout burst quota | Each supported Region: 20 per second | Yes | The maximum number of concurrent GetLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetLayout rate quota | Each supported Region: 6 per second | Yes | The maximum number of GetLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetTemplate burst quota | Each supported Region: 20 per second | Yes | The maximum number of concurrent GetTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| GetTemplate rate quota | Each supported Region: 6 per second | Yes | The maximum number of GetTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| Layouts per domain | Each supported Region: 100 | Yes | The maximum number of layouts in a Cases domain |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListCaseRules burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent ListCaseRules API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListCaseRules rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListCaseRules API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListCasesForContact burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent ListCasesForContact API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListCasesForContact rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListCasesForContact API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListDomains burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent ListDomains API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListDomains rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListDomains API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListFieldOptions burst quota | Each supported Region: 15 per second | Yes | The maximum number of concurrent ListFieldOptions API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListFieldOptions rate quota | Each supported Region: 6 per second | Yes | The maximum number of ListFieldOptions API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListFields burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent ListFields API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListFields rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListFields API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListLayouts burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent ListLayouts API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListLayouts rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListLayouts API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListTagsForResource burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent ListTagsForResource API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListTagsForResource rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListTagsForResource API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| ListTemplates burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent ListTemplates API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListTemplates rate quota | Each supported Region: 2 per second | Yes | The maximum number of ListTemplates API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| Parent field values per field options case rule | Each supported Region: 50 | Yes | The maximum number of parent field values per field options case rule |
| PutCaseEventConfiguration burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent PutCaseEventConfiguration API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| PutCaseEventConfiguration rate quota | Each supported Region: 2 per second | Yes | The maximum number of PutCaseEventConfiguration API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| Related items per case | Each supported Region: 200 | Yes | The maximum number of related items that can be attached to a case |
| SearchAllRelatedItems burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent SearchAllRelatedItems API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SearchAllRelatedItems rate quota | Each supported Region: 2 per second | Yes | The maximum number of SearchAllRelatedItems API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| SearchCases burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent SearchCases API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| SearchCases rate quota | Each supported Region: 2 per second | Yes | The maximum number of SearchCases API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SearchRelatedItems burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent SearchRelatedItems API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| SearchRelatedItems rate quota | Each supported Region: 2 per second | Yes | The maximum number of SearchRelatedItems API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| TagResource burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent TagResource API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| TagResource rate quota | Each supported Region: 2 per second | Yes | The maximum number of TagResource API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| Templates per domain | Each supported Region: 100 | Yes | The maximum number of templates in a Cases domain |
| UntagResource burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent UntagResource API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UntagResource rate quota | Each supported Region: 2 per second | Yes | The maximum number of UntagResource API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateCase burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent UpdateCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateCase rate quota | Each supported Region: 2 per second | Yes | The maximum number of UpdateCase API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateCaseRule burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent UpdateCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateCaseRule rate quota | Each supported Region: 2 per second | Yes | The maximum number of UpdateCaseRule API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateField burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent UpdateField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateField rate quota | Each supported Region: 2 per second | Yes | The maximum number of UpdateField API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateLayout burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent UpdateLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateLayout rate quota | Each supported Region: 2 per second | Yes | The maximum number of UpdateLayout API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateRelatedItem burst quota | Each supported Region: 10 per second | Yes | The maximum number of concurrent UpdateRelatedItem API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateRelatedItem rate quota | Each supported Region: 2 per second | Yes | The maximum number of UpdateRelatedItem API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |
| UpdateTemplate burst quota | Each supported Region: 5 per second | Yes | The maximum number of concurrent UpdateTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| UpdateTemplate rate quota | Each supported Region: 2 per second | Yes | The maximum number of UpdateTemplate API requests allowed per second. When you reach this quota, Amazon Connect Cases rejects requests for this operation for the remainder of the interval. |

For more information, see Connect Customer Service Quotas in the *Connect Customer Administrator Guide*.

## Amazon Q in Connect quotas

Amazon Q in Connect has the following quotas.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Assistant association | Each supported Region: 1 | No | The maximum number of assistant associations you can create for this assistant in the current Region. |
| Connect AI agent - assistant count | Each supported Region: 5 | No | The maximum number of Connect AI agent assistants you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Connect AI agent - knowledge base count | Each supported Region: 10 | No | The maximum number of Connect AI agent knowledge bases you can create in this account in the current Region. |
| Content per knowledge base | Each supported Region: 5,000 | No | The maximum pieces of content that you can add to a knowledge base in this Region. |
| Content size | Each supported Region: 1 Megabytes | No | The maximum size, in MB, of a piece of content that you can add to a knowledge base in this Region. |
| Maximum number of attachments per email message template | Each supported Region: 10 | No | The maximum number of attachments per email message template. After this limit is reached, Connect AI agent rejects requests to add more attachments to the email message template. |
| Maximum number of characters in a Baidu push message template | Each supported Region: 4,000 | No | The maximum number of characters in a Baidu push message template. |
| Maximum number of characters in a FCM push message template | Each supported Region: 4,000 | No | The maximum number of characters in a FCM push message template. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of characters in an ADM push message template | Each supported Region: 6,000 | No | The maximum number of characters in an ADM push message template. |
| Maximum number of characters in an APNs push message template | Each supported Region: 4,000 | No | The maximum number of characters in an APNs push message template. |
| Maximum number of characters in an SMS message template | Each supported Region: 800 | No | The maximum number of characters in an SMS message template. |
| Maximum number of characters in an email message template | Each supported Region: 5,000,000 | No | The maximum number of characters in an email message template. |
| Maximum number of message templates per knowledge base | Each supported Region: 200 | No | The maximum number of message templates per knowledge base. |
| Maximum number of versions per message template | Each supported Region: 20 | No | The maximum number of versions per message template. |
| Maximum size per attachment in an email message template | Each supported Region: 1 | No | The maximum size (in MB) of an attachment per email message template. Connect AI agent rejects requests to add larger attachments. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for ActivateMessageTemplate | Each supported Region: 3 | No | The maximum number of concurrent ActivateMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for CreateAssistant | Each supported Region: 10 | No | The maximum number of concurrent CreateAssistant API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for CreateAssistantAssociation | Each supported Region: 10 | No | The maximum number of concurrent CreateAssistantAssociation API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for CreateContent | Each supported Region: 10 | No | The maximum number of concurrent CreateContent API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for CreateKnowledgeBase | Each supported Region: 10 | No | The maximum number of concurrent CreateKnowledgeBase API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for CreateMessageTemplate | Each supported Region: 3 | No | The maximum number of concurrent CreateMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for CreateMessageTemplateAttachment | Each supported Region: 3 | No | The maximum number of concurrent CreateMessageTemplateAttachment API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for CreateMessageTemplateVersion | Each supported Region: 3 | No | The maximum number of concurrent CreateMessageTemplateVersion API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for CreateQuickResponse | Each supported Region: 10 | No | The maximum number of concurrent CreateQuickResponse API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for CreateSession | Each supported Region: 10 | No | The maximum number of concurrent CreateSession API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for DeactivateMessageTemplate | Each supported Region: 3 | No | The maximum number of concurrent DeactivateMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for DeleteAssistant | Each supported Region: 10 | No | The maximum number of concurrent DeleteAssistant API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for DeleteAssistantAssociation | Each supported Region: 10 | No | The maximum number of concurrent DeleteAssistantAssociation API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for DeleteContent | Each supported Region: 10 | No | The maximum number of concurrent DeleteContent API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for DeleteImportJob | Each supported Region: 10 | No | The maximum number of concurrent DeleteImportJob API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for DeleteKnowledgeBase | Each supported Region: 10 | No | The maximum number of concurrent DeleteKnowledgeBase API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for DeleteMessageTemplate | Each supported Region: 10 | No | The maximum number of concurrent DeleteMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for DeleteMessageTemplateAttachment | Each supported Region: 10 | No | The maximum number of concurrent DeleteMessageTemplateAttachment API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for DeleteQuickResponse | Each supported Region: 20 | No | The maximum number of concurrent DeleteQuickResponse API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetAssistant | Each supported Region: 10 | No | The maximum number of concurrent GetAssistant API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetAssistantAssociation | Each supported Region: 10 | No | The maximum number of concurrent GetAssistantAssociation API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for GetContent | Each supported Region: 10 | No | The maximum number of concurrent GetContent API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetContentSummary | Each supported Region: 10 | No | The maximum number of concurrent GetContentSummary API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetImportJob | Each supported Region: 10 | No | The maximum number of concurrent GetImportJob API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of API requests for GetKnowledgeBase | Each supported Region: 10 | No | The maximum number of concurrent GetKnowledgeBase API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetMessageTemplate | Each supported Region: 10 | No | The maximum number of concurrent GetMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetQuickResponse | Each supported Region: 10 | No | The maximum number of concurrent GetQuickResponse API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for GetRecommendations | Each supported Region: 10 | No | The maximum number of concurrent GetRecommendations API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for GetSession | Each supported Region: 10 | No | The maximum number of concurrent GetSession API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for ListAssistantAssociations | Each supported Region: 10 | No | The maximum number of concurrent ListAssistantAssociations API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for ListAssistants | Each supported Region: 10 | No | The maximum number of concurrent ListAssistants API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for ListImportJobs | Each supported Region: 10 | No | The maximum number of concurrent ListImportJobs API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for ListKnowledgeBases | Each supported Region: 10 | No | The maximum number of concurrent ListKnowledgeBases API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for ListMessageTemplateVersions | Each supported Region: 10 | No | The maximum number of concurrent ListMessageTemplateVersions API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for ListMessageTemplates | Each supported Region: 10 | No | The maximum number of concurrent ListMessageTemplates API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for ListQuickResponses | Each supported Region: 10 | No | The maximum number of concurrent ListQuickResponses API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for ListTagsForResource | Each supported Region: 10 | No | The maximum number of concurrent ListTagsForResource API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for NotifyRecommendationsReceived | Each supported Region: 10 | No | The maximum number of concurrent NotifyRecommendationsReceived API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for QueryAssistant | Each supported Region: 10 | No | The maximum number of concurrent QueryAssistant API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for RemoveKnowledgeBaseTemplateUri | Each supported Region: 10 | No | The maximum number of concurrent RemoveKnowledgeBaseTemplateUri API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for RenderMessageTemplate | Each supported Region: 10 | No | The maximum number of concurrent RenderMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for SearchContent | Each supported Region: 10 | No | The maximum number of concurrent SearchContent API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for SearchMessageTemplates | Each supported Region: 10 | No | The maximum number of concurrent SearchMessageTemplates API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for SearchQuickResponses | Each supported Region: 20 | No | The maximum number of concurrent SearchQuickResponses API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for SearchSessions | Each supported Region: 10 | No | The maximum number of concurrent SearchSessions API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for StartContentUpload | Each supported Region: 10 | No | The maximum number of concurrent StartContentUpload API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for StartImportJob | Each supported Region: 10 | No | The maximum number of concurrent StartImportJob API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for TagResource | Each supported Region: 10 | No | The maximum number of concurrent TagResource API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for UntagResource | Each supported Region: 10 | No | The maximum number of concurrent UntagResource API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for UpdateContent | Each supported Region: 10 | No | The maximum number of concurrent UpdateContent API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for UpdateKnowledgeBaseTemplateUri | Each supported Region: 10 | No | The maximum number of concurrent UpdateKnowledgeBaseTemplateUri API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of API requests for UpdateMessageTemplate | Each supported Region: 3 | No | The maximum number of concurrent UpdateMessageTemplate API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for UpdateMessageTemplateMetadata | Each supported Region: 3 | No | The maximum number of concurrent UpdateMessageTemplateMetadata API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |
| Rate of API requests for UpdateQuickResponse | Each supported Region: 10 | No | The maximum number of concurrent UpdateQuickResponse API requests allowed per second. When you reach this quota, Connect AI agent rejects requests for this operation for the remainder of the interval. |

# Amazon Connect Health endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

For information about Amazon Connect Health endpoints, see Supported Regions in the *Amazon Connect Health User Guide*.

## Service quotas

For information about Amazon Connect Health quotas, see Quotas for Amazon Connect Health in the *Amazon Connect Health User Guide*.

# AWS Control Tower endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service Endpoints

AWS Control Tower has the following endpoints:

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | controltower.us-east-2.amazonaws.com | HTTPS | |
| | | controltower-fips.us-east-2.api.aws | HTTPS | |
| | | controltower-fips.us-east-2.amazonaws.com | HTTPS | |
| | | controltower.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | controltower.us-east-1.amazonaws.com | HTTPS | |
| | | controltower-fips.us-east-1.api.aws | HTTPS | |
| | | controltower-fips.us-east-1.amazonaws.com | HTTPS | |
| | | controltower.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | controltower.us-west-1.amazonaws.com | HTTPS | |
| | | controltower-fips.us-west-1.api.aws | HTTPS | |
| | | controltower-fips.us-west-1.amazonaws.com | HTTPS | |
| | | controltower.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | controltower.us-west-2.amazonaws.com | HTTPS | |
| | | controltower-fips.us-west-2.api.aws | HTTPS | |
| | | controltower-fips.us-west-2.amazonaws.com | HTTPS | |
| | | controltower.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | controltower.af-south-1.amazonaws.com | HTTPS | |
| | | controltower.af-south-1.api.aws | HTTPS | |
| Asia Pacific | ap-east-1 | controltower.ap-east-1.amazonaws.com | HTTPS | |
| | | controltower.ap-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | controltower.ap-south-2.amazonaws.com<br>controltower.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | controltower.ap-southeast-3.amazonaws.com<br>controltower.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | controltower.ap-southeast-5.amazonaws.com<br>controltower.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | controltower.ap-southeast-4.amazonaws.com<br>controltower.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | controltower.ap-south-1.amazonaws.com<br>controltower.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | controltower.ap-southeast-6.amazonaws.com<br>controltower.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | controltower.ap-northeast-3.amazonaws.com<br>controltower.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | controltower.ap-northeast-2.amazonaws.com<br>controltower.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | controltower.ap-southeast-1.amazonaws.com<br>controltower.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | controltower.ap-southeast-2.amazonaws.com<br>controltower.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | controltower.ap-east-2.amazonaws.com<br>controltower.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | controltower.ap-southeast-7.amazonaws.com<br>controltower.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | controltower.ap-northeast-1.amazonaws.com<br>controltower.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | controltower.ca-central-1.amazonaws.com<br>controltower-fips.ca-central-1.api.aws<br>controltower-fips.ca-central-1.amazonaws.com<br>controltower.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | controltower.ca-west-1.amazonaws.com | HTTPS | |
| | | controltower-fips.ca-west-1.api.aws | HTTPS | |
| | | controltower-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | controltower.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | controltower.eu-central-1.amazonaws.com | HTTPS | |
| | | controltower.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | controltower.eu-west-1.amazonaws.com | HTTPS | |
| | | controltower.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | controltower.eu-west-2.amazonaws.com | HTTPS | |
| | | controltower.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | controltower.eu-south-1.amazonaws.com | HTTPS | |
| | | controltower.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | controltower.eu-west-3.amazonaws.com | HTTPS | |
| | | controltower.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | controltower.eu-south-2.amazonaws.com | HTTPS | |
| | | controltower.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | controltower.eu-north-1.amazonaws.com | HTTPS | |
| | | controltower.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | controltower.eu-central-2.amazonaws.com | HTTPS | |
| | | controltower.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | controltower.il-central-1.amazonaws.com<br>controltower.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | controltower.mx-central-1.amazonaws.com<br>controltower.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | controltower.me-south-1.amazonaws.com<br>controltower.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | controltower.me-central-1.amazonaws.com<br>controltower.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | controltower.sa-east-1.amazonaws.com<br>controltower.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | controltower.us-gov-east-1.amazonaws.com<br>controltower-fips.us-gov-east-1.api.aws<br>controltower-fips.us-gov-east-1.amazonaws.com<br>controltower.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | controltower.us-gov-west-1.amazonaws.com<br>controltower-fips.us-gov-west-1.api.aws<br>controltower-fips.us-gov-west-1.amazonaws.com<br>controltower.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

# Service Quotas

AWS Control Tower has the following quotas.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent account operations quota | Each supported Region: 5 | Yes | The maximum number of concurrent account operations that can be performed at the same time. Maximum: 10 |
| Concurrent organization units (OUs) operations quota | Each supported Region: 1 | No | The maximum number of concurrent OU-related operations that can be performed at the same time. |
| Number of accounts in a single OU quota | Each supported Region: 1,000 | No | The maximum number of AWS Control Tower managed accounts that can be present in one OU. If you add accounts beyond this limit, the OU registration process in AWS Control Tower cannot be performed. To learn more about the number of accounts per OU, review Regions and Stackset Limitations in the AWS Control Tower documentation. |

# AWS Data Exchange endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | dataexchange.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | dataexchange.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | dataexchange.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | dataexchange.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | dataexchange.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific | ap-southeast-1 | dataexchange.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Singapore) | | | | |
| Asia Pacific (Sydney) | ap-southeast-2 | dataexchange.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | dataexchange.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | dataexchange.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | dataexchange.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | dataexchange.eu-west-2.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AWS Lake Formation data permission assets per revision | Each supported Region: 20 | Yes | The maximum number of AWS Lake Formation data permission assets that a single revision can contain. |
| Active and pending data grants | Each supported Region: 25,000 | No | The maximum number of active and pending data |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | grants from a single data owner. |
| Amazon API Gateway API assets per revision | Each supported Region: 20 | Yes | The maximum number of Amazon API Gateway API assets that a single revision can contain. |
| Amazon Redshift datashare assets per import job from Redshift | Each supported Region: 10 | No | The maximum number of Amazon Redshift datashare assets you can import from Redshift in a single job. |
| Amazon Redshift datashare assets per revision | Each supported Region: 20 | Yes | The maximum number of Amazon Redshift datashare assets that a single revision can contain. |
| Amazon S3 data access assets per revision | Each supported Region: 5 | Yes | The maximum number of Amazon S3 data access assets that a single revision can contain. |
| Asset per export job from Amazon S3 | Each supported Region: 100 | No | The maximum number of assets you can export to Amazon S3 in a single job. |
| Asset size in GB | Each supported Region: 100 Gigabytes | No | The maximum size in GB of a single asset. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Assets per import job from Amazon S3 | Each supported Region: 100 | No | The maximum number of assets you can import from Amazon S3 in a single job. |
| Assets per revision | Each supported Region: 10,000 | Yes | The maximum number of assets that a single revision can contain. |
| Auto export event actions per data set | Each supported Region: 5 | Yes | The maximum number of auto export event actions per data set. |
| Bring-Your-Own-Subscription offers per account | Each supported Region: 10 | Yes | The maximum number of BYOS (Bring-Your-Own-Subscription) offers that a single account can create. |
| Concurrent in progress jobs to create Amazon S3 data access assets from S3 buckets | Each supported Region: 5 | No | The maximum number of concurrently running jobs (i.e jobs IN_PROGRESS) to create Amazon S3 data access assets from S3 buckets. |
| Concurrent in progress jobs to create data grants | Each supported Region: 5 | No | The maximum number of concurrently running jobs (i.e jobs IN_PROGRESS) to create data grants. |
| Concurrent in progress jobs to delete data grants | Each supported Region: 5 | No | The maximum number of concurrently running jobs (i.e jobs IN_PROGRESS) to delete data grants. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent in progress jobs to export assets to Amazon S3 | Each supported Region: 10 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to export assets to Amazon S3. |
| Concurrent in progress jobs to export assets to a signed URL | Each supported Region: 10 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to export assets to a signed URL. |
| Concurrent in progress jobs to export revisions to Amazon S3 | Each supported Region: 5 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to export revisions to Amazon S3. |
| Concurrent in progress jobs to import assets from Amazon API Gateway | Each supported Region: 10 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to import assets from Amazon API Gateway. |
| Concurrent in progress jobs to import assets from Amazon Redshift datashares | Each supported Region: 10 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to import Amazon Redshift datashares. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent in progress jobs to import assets from Amazon S3 | Each supported Region: 10 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to import assets from Amazon S3. |
| Concurrent in progress jobs to import assets from a signed URL | Each supported Region: 10 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to import assets from a signed URL. |
| Concurrent in progress jobs to import assets from an AWS Lake Formation tag policy | Each supported Region: 5 | No | The maximum number of concurrently running jobs (i.e. jobs with IN_PROGRESS state) to import assets from an AWS Lake Formation tag policy. |
| Data dictionaries per product | Each supported Region: 50 | Yes | The maximum number of data dictionaries per product. |
| Data dictionary file size in MB | Each supported Region: 1 Megabytes | Yes | The maximum size (in MB) of a data dictionary. |
| Data sets per account | Each supported Region: 3,000 | Yes | The maximum number of data sets per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Data sets per product | Each supported Region: 25 | Yes | The maximum number of data sets that a single product can contain. |
| Event actions per account | Each supported Region: 50 | Yes | The maximum number of event actions per account. |
| Pending data grants per consumer | Each supported Region: 20 | Yes | The maximum number of pending data grants from a single data owner to a single consumer account. |
| Private offers per account | Each supported Region: 25 | Yes | The maximum number of custom offers that a single account can create. |
| Products per account | Each supported Region: 50 | Yes | The maximum number of published products per account. |
| Products per data set | Each supported Region: 100 | Yes | The maximum number of products that can contain a given data set. |
| Revisions per AWS Lake Formation data permission data set | Each supported Region: 20 | Yes | The maximum number of revisions that a single AWS Lake Formation data permission data set can contain. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Revisions per Amazon API Gateway API data set | Each supported Region: 20 | Yes | The maximum number of revisions that a single Amazon API Gateway API data set can contain. |
| Revisions per Amazon Redshift datashare data set | Each supported Region: 20 | Yes | The maximum number of revisions that a single Amazon Redshift datashare data set can contain. |
| Revisions per Amazon S3 data access data set | Each supported Region: 5 | Yes | The maximum number of revisions that a single Amazon S3 data access data set can contain. |
| Revisions per addRevisions change set | Each supported Region: 5 | No | The maximum number of revisions that can be published to a product in a single AWS Marketplace Catalog API ChangeSet of type addRevisions. |
| Revisions per data set | Each supported Region: 10,000 | Yes | The maximum number of revisions that a single data set can contain. |
| Sample file size in MB | Each supported Region: 50 Megabytes | Yes | The maximum size (in MB) of a sample. |
| Samples per product | Each supported Region: 50 | Yes | The maximum number of samples per product. |

For more information, see AWS Data Exchange quotas in the *AWS Data Exchange User Guide*.

# Amazon DataZone endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | datazone.us-east-2.api.aws | HTTPS |
| | | datazone-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | datazone.us-east-1.api.aws | HTTPS |
| | | datazone-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | datazone.us-west-2.api.aws | HTTPS |
| | | datazone-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | datazone.ap-east-1.api.aws | HTTPS |
| Asia Pacific | ap-southeast-5 | datazone.ap-southeast-5.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Malaysia) | | | | |
| Asia Pacific (Mumbai) | ap-south-1 | datazone.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | datazone.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | datazone.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | datazone.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | datazone.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | datazone.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | datazone.ca-central-1.api.aws<br>datazone-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | datazone.eu-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | datazone.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | datazone.eu-west-2.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | datazone.eu-west-3.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | datazone.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | datazone.eu-central-2.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | datazone.sa-east-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Asset Types | Each supported Region: 1,000 | Yes | The maximum number of asset types that can be created in an Amazon DataZone domain. |
| Assets | Each supported Region: 1,000,000 | Yes | The maximum number of assets that can be |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | created in an Amazon DataZone domain. |
| Business Glossary Terms | Each supported Region: 10,000 | Yes | The maximum number of business glossary terms that can be created in an Amazon DataZone domain. |
| Data Products | Each supported Region: 500,000 | Yes | The maximum number of data products that can be created in an Amazon DataZone domain. |
| Data source runs/data source/day | Each supported Region: 25 | Yes | The maximum data source runs per data source per day. |
| Environments in a domain | Each supported Region: 1,000 | Yes | The maximum number of Amazon DataZone environments in an Amazon DataZone domain. |
| Git connections per project | Each supported Region: 5 | Yes | The maximum number of connections of type GIT that can be created in a project |
| Glossaries | Each supported Region: 1,000 | Yes | The maximum number of business glossaries that can be created in an Amazon DataZone domain. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Compute environments running on sc.t3.large instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Compute environments running on sc.t3.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.12x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.18xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.18x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.24x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.9xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.9xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.c5.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c5.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c5.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.12 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.16 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.24 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.32 xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6i.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.16xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.24xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.32xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.c6id.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c6id.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.12xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.16xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.24xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.48xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.c7i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.c7i.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g4dn.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g4dn.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g4dn.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g4dn.16xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g4dn.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g4dn.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g4dn.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g4dn.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g4dn.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g4dn.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g4dn.xlarge instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g4dn.xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.g5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.12x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.g5.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.16x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.g5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.24x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.g5.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g5.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.48x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.g5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.g5.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.g5.xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.g5.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.12x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.16x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.24x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.m5.large instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5.xlarge instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.12 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.16 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.24 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.2x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.4x large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.8x large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.large instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m5d.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m5d.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.12 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.16 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.24 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.2x large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.32xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.large instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6i.xlarge instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6i.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.16xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.24xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.32xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.large instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m6id.xlarge instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m6id.xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.16xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.24xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.48xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.large instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.m7i.xlarge instances | Each supported Region: 2 | Yes | SageMaker Unified Studio Notebook Runs running on sc.m7i.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.16xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.24xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.r5.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.r5.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r5.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.12 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.16 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.24 xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.2xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.32xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6i.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.16xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.24xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.32xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.8xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.r6id.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r6id.xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.12xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.16xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.24xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.48xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.4xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.r7i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebook Runs running on sc.r7i.xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.t3.2xlarge instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Notebook Runs running on sc.t3.2xlarge instances |
| SageMaker Unified Studio Notebook Runs running on sc.t3.large instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Notebook Runs running on sc.t3.large instances |
| SageMaker Unified Studio Notebook Runs running on sc.t3.medium instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Notebook Runs running on sc.t3.medium instances |
| SageMaker Unified Studio Notebook Runs running on sc.t3.xlarge instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Notebook Runs running on sc.t3.xlarge instances |
| SageMaker Unified Studio Notebooks | Each supported Region: 5,000 | Yes | The maximum number of notebooks that can be created in a domain. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.c5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.12x large instances |
| SageMaker Unified Studio Notebooks running on sc.c5.18xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.18x large instances |
| SageMaker Unified Studio Notebooks running on sc.c5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.24x large instances |
| SageMaker Unified Studio Notebooks running on sc.c5.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c5.9xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.9xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c5.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.c5.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c5.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.32xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.c6i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.large instances |
| SageMaker Unified Studio Notebooks running on sc.c6i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c6i.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.2xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.c6id.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.32xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.large instances |
| SageMaker Unified Studio Notebooks running on sc.c6id.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c6id.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.16xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.c7i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.large instances |
| SageMaker Unified Studio Notebooks running on sc.c7i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.c7i.xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.g4dn.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g4dn.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g4dn.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g4dn.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g4dn.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g4dn.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g4dn.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g4dn.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g4dn.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g4dn.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g4dn.xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g4dn.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.12xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.g5.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g5.xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g5.xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.g6.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.12x large instances |
| SageMaker Unified Studio Notebooks running on sc.g6.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.16x large instances |
| SageMaker Unified Studio Notebooks running on sc.g6.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.24x large instances |
| SageMaker Unified Studio Notebooks running on sc.g6.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.48x large instances |
| SageMaker Unified Studio Notebooks running on sc.g6.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.g6.xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.g6e.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.8x large instances |
| SageMaker Unified Studio Notebooks running on sc.g6e.xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.g6e.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.12x large instances |
| SageMaker Unified Studio Notebooks running on sc.m5.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.16x large instances |
| SageMaker Unified Studio Notebooks running on sc.m5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.24x large instances |
| SageMaker Unified Studio Notebooks running on sc.m5.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.m5.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.large instances |
| SageMaker Unified Studio Notebooks running on sc.m5.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m5.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.12 xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.16 xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.24 xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.2x large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.m5d.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.large instances |
| SageMaker Unified Studio Notebooks running on sc.m5d.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m5d.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.24xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.m6i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.32xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.large instances |
| SageMaker Unified Studio Notebooks running on sc.m6i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m6i.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.12xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.m6id.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.32xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m6id.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.m6id.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m6id.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.2xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.m7i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.8x large instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.large instances |
| SageMaker Unified Studio Notebooks running on sc.m7i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.m7i.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p3.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p3.16x large instances |
| SageMaker Unified Studio Notebooks running on sc.p3.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p3.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p3.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p3.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p3dn.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p3dn.2 4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.p4d.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p4d.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p4de.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p4de.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p5.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p5.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p5en.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p5en.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.p6-b200.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.p6-b200.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r5.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r5.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.16xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.r5.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r5.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r5.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r5.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r5.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.large instances |
| SageMaker Unified Studio Notebooks running on sc.r5.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r5.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.12 xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.r6i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.32xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.r6i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r6i.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.32xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.32xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.4xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.r6id.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.large instances |
| SageMaker Unified Studio Notebooks running on sc.r6id.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r6id.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.12xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.12xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.16xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.16xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.24xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.24xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.2xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.2xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.r7i.48xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.48xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.4xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.4xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.8xlarge instances | Each supported Region: 0 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.8xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.large instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.large instances |
| SageMaker Unified Studio Notebooks running on sc.r7i.xlarge instances | Each supported Region: 1 | Yes | SageMaker Unified Studio Notebooks running on sc.r7i.xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.t3.2xlarge instances | Each supported Region: 6 | Yes | SageMaker Unified Studio Notebooks running on sc.t3.2xlarge instances |
| SageMaker Unified Studio Notebooks running on sc.t3.large instances | Each supported Region: 6 | Yes | SageMaker Unified Studio Notebooks running on sc.t3.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SageMaker Unified Studio Notebooks running on sc.t3.medium instances | Each supported Region: 6 | Yes | SageMaker Unified Studio Notebooks running on sc.t3.medium instances |
| SageMaker Unified Studio Notebooks running on sc.t3.xlarge instances | Each supported Region: 6 | Yes | SageMaker Unified Studio Notebooks running on sc.t3.xlarge instances |
| SageMaker Unified Studio Querybook Runs running on sc.t3.medium instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Querybook Runs running on sc.t3.medium instances |
| SageMaker Unified Studio Querybooks running on sc.t3.medium instances | Each supported Region: 5 | Yes | SageMaker Unified Studio Querybooks running on sc.t3.medium instances |
| SageMaker Unified Studio agent requests per month | Each supported Region: 1,000 | No | The maximum number of requests that can be made to the SageMaker Unified Studio agent per month. |
| The maximum number of concurrent notebook runs per account. | Each supported Region: 15 | Yes | The maximum number of concurrent notebook runs per account. |
| Total EBS volume size in GB across all SageMaker Unified Studio Compute environments | Each supported Region: 500 | Yes | Total EBS volume size in GB across all SageMaker Unified Studio Compute environments |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Total EBS volume size in GB across all SageMaker Unified Studio Notebooks | Each supported Region: 500 | Yes | Total EBS volume size in GB across all SageMaker Unified Studio Notebooks |

For details, see Quotas for Amazon DataZone.

# Amazon Data Lifecycle Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | dlm.us-east-2.amazonaws.com | HTTPS |
| | | dlm.us-east-2.api.aws | HTTPS |
| | | dlm-fips.us-east-2.amazonaws.com | HTTPS |
| | | dlm-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | dlm.us-east-1.amazonaws.com | HTTPS |
| | | dlm.us-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | dlm-fips.us-east-1.api.aws | HTTPS |
| | | dlm-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | dlm.us-west-1.amazonaws.com | HTTPS |
| | | dlm.us-west-1.api.aws | HTTPS |
| | | dlm-fips.us-west-1.api.aws | HTTPS |
| | | dlm-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | dlm.us-west-2.amazonaws.com | HTTPS |
| | | dlm-fips.us-west-2.amazonaws.com | HTTPS |
| | | dlm.us-west-2.api.aws | HTTPS |
| | | dlm-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | dlm.af-south-1.amazonaws.com | HTTPS |
| | | dlm.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | dlm.ap-east-1.amazonaws.com | HTTPS |
| | | dlm.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | dlm.ap-south-2.amazonaws.com | HTTPS |
| | | dlm.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | dlm.ap-southeast-3.amazonaws.com | HTTPS |
| | | dlm.ap-southeast-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | dlm.ap-southeast-5.amazonaws.com<br>dlm.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | dlm.ap-southeast-4.amazonaws.com<br>dlm.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | dlm.ap-south-1.amazonaws.com<br>dlm.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | dlm.ap-southeast-6.amazonaws.com<br>dlm.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | dlm.ap-northeast-3.amazonaws.com<br>dlm.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | dlm.ap-northeast-2.amazonaws.com<br>dlm.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | dlm.ap-southeast-1.amazonaws.com<br>dlm.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | dlm.ap-southeast-2.amazonaws.com<br>dlm.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | dlm.ap-east-2.amazonaws.com | HTTPS |
| | | dlm.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | dlm.ap-southeast-7.amazonaws.com | HTTPS |
| | | dlm.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | dlm.ap-northeast-1.amazonaws.com | HTTPS |
| | | dlm.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | dlm.ca-central-1.amazonaws.com | HTTPS |
| | | dlm.ca-central-1.api.aws | HTTPS |
| | | dlm-fips.ca-central-1.api.aws | HTTPS |
| | | dlm-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | dlm.ca-west-1.amazonaws.com | HTTPS |
| | | dlm.ca-west-1.api.aws | HTTPS |
| | | dlm-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | dlm.eu-central-1.amazonaws.com | HTTPS |
| | | dlm.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | dlm.eu-west-1.amazonaws.com | HTTPS |
| | | dlm.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | dlm.eu-west-2.amazonaws.com | HTTPS |
| | | dlm.eu-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | dlm.eu-south-1.amazonaws.com | HTTPS | |
| | | dlm.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | dlm.eu-west-3.amazonaws.com | HTTPS | |
| | | dlm.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | dlm.eu-south-2.amazonaws.com | HTTPS | |
| | | dlm.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | dlm.eu-north-1.amazonaws.com | HTTPS | |
| | | dlm.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | dlm.eu-central-2.amazonaws.com | HTTPS | |
| | | dlm.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | dlm.il-central-1.amazonaws.com | HTTPS | |
| | | dlm.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | dlm.mx-central-1.amazonaws.com | HTTPS | |
| | | dlm.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | dlm.me-south-1.amazonaws.com | HTTPS | |
| | | dlm.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | dlm.me-central-1.amazonaws.com | HTTPS | |
| | | dlm.me-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | dlm.sa-east-1.amazonaws.com<br>dlm.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | dlm.us-gov-east-1.amazonaws.com<br>dlm.us-gov-east-1.api.aws<br>dlm-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | dlm.us-gov-west-1.amazonaws.com<br>dlm.us-gov-west-1.api.aws<br>dlm-fips.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Policies per Region | Each supported Region: 100 | Yes | The maximum number of policies per Region. |
| Target accounts per sharing rule | Each supported Region: 50 | Yes | The maximum number of target accounts per sharing rule. |

# AWS Data Pipeline endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | datapipeline.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | datapipeline.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | datapipeline.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | datapipeline.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | datapipeline.eu-west-1.amazonaws.com | HTTPS | |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Minimum delay between retry attempts in minutes | Each supported Region: 2 | No | The minimum delay between retry attempts in minutes. |
| Minimum scheduling interval in minutes | Each supported Region: 15 | No | The minimum scheduling interval in minutes. |
| Number of EC2 instances per Ec2Resource object | Each supported Region: 1 | No | The maximum number of EC2 instances per Ec2Resource object. |
| Number of UTF8 bytes per field | Each supported Region: 10,240 | No | The maximum number of UTF8 bytes per field. |
| Number of UTF8 bytes per field name or identifier | Each supported Region: 256 | No | The maximum number of UTF8 bytes per field name or identifier. |
| Number of UTF8 bytes per object | Each supported Region: 15,360 | No | The maximum number of UTF8 bytes per object (including field names). |
| Number of active instances per object | Each supported Region: 5 | Yes | The maximum number of active instances per object. |
| Number of fields per object | Each supported Region: 50 | No | The maximum number of fields per object. |
| Number of objects per pipeline | Each supported Region: 100 | Yes | The maximum number of objects that you can define per pipeline. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Number of pipelines you can create | Each supported Region: 100 | Yes | The maximum number of pipelines that you can create. |
| Number of roll-ups into a single object | Each supported Region: 32 | No | The maximum number of roll-ups into a single object. |
| Rate of creation of an instance from an object | Each supported Region: 1 per 5 minutes | No | The rate of creation of an instance from an object per 5 minutes. |
| Retries of a pipeline activity per task | Each supported Region: 5 | No | The maximum number of retries of a pipeline activity per task. |

For more information, see AWS Data Pipeline Quotas in the *AWS Data Pipeline Developer Guide*.

# AWS DataSync endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Topics

- DataSync transfers

## DataSync transfers

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | datasync.us-east-2.amazonaws.com | HTTPS |
| | | datasync-fips.us-east-2.api.aws | HTTPS |
| | | datasync.us-east-2.api.aws | HTTPS |
| | | datasync-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | datasync.us-east-1.amazonaws.com | HTTPS |
| | | datasync-fips.us-east-1.api.aws | HTTPS |
| | | datasync-fips.us-east-1.amazonaws.com | HTTPS |
| | | datasync.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | datasync.us-west-1.amazonaws.com | HTTPS |
| | | datasync.us-west-1.api.aws | HTTPS |
| | | datasync-fips.us-west-1.api.aws | HTTPS |
| | | datasync-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | datasync.us-west-2.amazonaws.com | HTTPS |
| | | datasync-fips.us-west-2.api.aws | HTTPS |
| | | datasync.us-west-2.api.aws | HTTPS |
| | | datasync-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | datasync.af-south-1.amazonaws.com | HTTPS |
| | | datasync.af-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | datasync.ap-east-1.amazonaws.com<br>datasync.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | datasync.ap-south-2.amazonaws.com<br>datasync.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | datasync.ap-southeast-3.amazonaws.com<br>datasync.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | datasync.ap-southeast-5.amazonaws.com<br>datasync.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | datasync.ap-southeast-4.amazonaws.com<br>datasync.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | datasync.ap-south-1.amazonaws.com<br>datasync.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | datasync.ap-southeast-6.amazonaws.com<br>datasync.ap-southeast-6.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | datasync.ap-northeast-3.amazonaws.com | HTTPS |
| | | datasync.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | datasync.ap-northeast-2.amazonaws.com | HTTPS |
| | | datasync.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | datasync.ap-southeast-1.amazonaws.com | HTTPS |
| | | datasync.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | datasync.ap-southeast-2.amazonaws.com | HTTPS |
| | | datasync.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | datasync.ap-east-2.amazonaws.com | HTTPS |
| | | datasync.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | datasync.ap-southeast-7.amazonaws.com | HTTPS |
| | | datasync.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | datasync.ap-northeast-1.amazonaws.com | HTTPS |
| | | datasync.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | datasync.ca-central-1.amazonaws.com | HTTPS |
| | | datasync.ca-central-1.api.aws | HTTPS |
| | | datasync-fips.ca-central-1.amazonaws.com | HTTPS |
| | | datasync-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | datasync.ca-west-1.amazonaws.com | HTTPS |
| | | datasync.ca-west-1.api.aws | HTTPS |
| | | datasync-fips.ca-west-1.amazonaws.com | HTTPS |
| | | datasync-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | datasync.eu-central-1.amazonaws.com | HTTPS |
| | | datasync.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | datasync.eu-west-1.amazonaws.com | HTTPS |
| | | datasync.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | datasync.eu-west-2.amazonaws.com | HTTPS |
| | | datasync.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | datasync.eu-south-1.amazonaws.com | HTTPS |
| | | datasync.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | datasync.eu-west-3.amazonaws.com | HTTPS |
| | | datasync.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | datasync.eu-south-2.amazonaws.com | HTTPS |
| | | datasync.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | datasync.eu-north-1.amazonaws.com | HTTPS |
| | | datasync.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | datasync.eu-central-2.amazonaws.com | HTTPS |
| | | datasync.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | datasync.il-central-1.amazonaws.com | HTTPS |
| | | datasync.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | datasync.mx-central-1.amazonaws.com | HTTPS |
| | | datasync.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | datasync.me-south-1.amazonaws.com | HTTPS |
| | | datasync.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | datasync.me-central-1.amazonaws.com | HTTPS |
| | | datasync.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | datasync.sa-east-1.amazonaws.com | HTTPS |
| | | datasync.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | datasync.us-gov-east-1.amazonaws.com | HTTPS |
| | | datasync.us-gov-east-1.api.aws | HTTPS |
| | | datasync-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | datasync-fips.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | datasync.us-gov-west-1.amazonaws.com | HTTPS | |
| | | datasync-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | datasync-fips.us-gov-west-1.api.aws | HTTPS | |
| | | datasync.us-gov-west-1.api.aws | HTTPS | |

# Service quotas

### Topics

- DataSync task quotas

## DataSync task quotas

The following table describes the quotas for DataSync tasks in a specific AWS account and AWS Region.

| Resource | Quota | Adjustable |
|---|---|---|
| Maximum number of tasks you can create | 100 | Yes |
| **(Enhanced mode tasks)** Maximum number of source and destination objects that DataSync can work with per task execution<br><br>For more information, see How DataSync transfers files, objects, and directories | Virtually unlimited | N/A |
| **(Basic mode tasks)** Maximum number of source and destination files, objects, and directories that DataSync can work with per task execution between self-managed storage or other clouds and AWS services | 50 million | Yes<br><br>**Tip**<br>Instead of |

| Resource | Quota | Adjustable |
| --- | --- | --- |
| For more information, see How DataSync transfers files, objects, and directories | **Important**<br>Remember the following about this quota:<br>• If you transfer Amazon S3 objects with prefixes, the prefixes are treated as directories and count towards the quota. For example, DataSync would considers3://bucket/foo/bar.txt as two directories (./ and ./foo/) and one object (bar.txt).<br>• If your task is transferring more than 20 million files, objects, or directories, make sure that you allocate a minimum of 64 GB of RAM to your DataSync agent. For more information, see agent requirements for transfers. | requesting an increase, you can create tasks that focus on specific directories using include and exclude filters. For more information, see Filtering the data transferred by AWS DataSync. |

| Resource | Quota | Adjustable |
| --- | --- | --- |
| **(Basic mode tasks)** Maximum number of source and destination files, objects, and directories that DataSync can work with per task execution between AWS storage services<br><br>For more information, see How DataSync transfers files, objects, and directories | 25 million<br><br>**Important**<br>If you transfer Amazon S3 objects with prefixes, the prefixes are treated as directories and count towards the quota. For example, DataSync would considers3://bucket/foo/bar.txt as two directories (./ and ./foo/) and one object (bar.txt). | Yes<br><br>**Tip**<br>Instead of requesting an increase, you can create tasks that focus on specific directories using include and exclude filters. For more information, see Filtering the data transferred by |

| Resource | Quota | Adjustable |
| --- | --- | --- |
| | | AWS DataSync. |
| Maximum throughput per task (for transfers that use a DataSync agent) | 10 Gbps | No |
| Maximum throughput per task (for transfers that don't use a DataSync agent) | 5 Gbps | No |
| Maximum number of characters you can include in a task filter | 102,400 characters<br>**Note**<br>If you're using the DataSync console, this limit includes all the characters combined in your include and exclude patterns. | No |
| Maximum number of queued executions for a single task | 50 | No |
| Maximum number of concurrent Enhanced mode task executions | 120 | No |
| Number of days a task execution's history is retained | 30 | No |

## AWS Database Migration Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | dms.us-east-2.amazonaws.com<br>dms-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | dms.us-east-1.amazonaws.com<br>dms-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | dms.us-west-1.amazonaws.com<br>dms-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | dms.us-west-2.amazonaws.com<br>dms-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | dms.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | dms.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-south-2 | dms.ap-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | dms.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | dms.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | dms.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | dms.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | dms.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | dms.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | dms.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | dms.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | dms.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | dms.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | dms.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | dms.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | dms.ca-central-1.amazonaws.com<br>dms-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | dms.ca-west-1.amazonaws.com<br>dms-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | dms.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | dms.eu-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | dms.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | dms.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | dms.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | dms.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | dms.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | dms.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | dms.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | dms.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | dms.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | dms.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | dms.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | dms.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | dms.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Certificate count | Each supported Region: 100 | Yes | The maximum number of certificates allowed in this account in the current Region. |
| Endpoint count | Each supported Region: 1,000 | Yes | The maximum number of endpoints allowed in this account in the current Region. |
| Endpoints per instance | Each supported Region: 100 | Yes | The maximum number of endpoints allowed per replication instance. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Event subscriptions | Each supported Region: 60 | Yes | The maximum number of AWS DMS events that you can subscribe to. |
| Number of DMS Fleet Advisor collector instances | Each supported Region: 10 | No | Maximum number of created instances of DMS Fleet Advisor collector |
| Number of Data Migrations | Each supported Region: 1,000 | Yes | Maximum number of data migrations |
| Number of Data Migrations per Migration Project | Each supported Region: 10 | No | Maximum number of data migrations allowed per migration project |
| Number of Data Providers | Each supported Region: 1,000 | Yes | Maximum number of data providers |
| Number of Instance Profiles | Each supported Region: 60 | Yes | Maximum number of instance profiles |
| Number of Migration Projects | Each supported Region: 50 | Yes | Maximum number of migration projects |
| Number of data files DMS Fleet Advisor can send per hour | Each supported Region: 500 per hour | No | Maximum number of files that DMS Fleet Advisor collector can send per hour |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of database objects DMS Fleet Advisor can process | Each supported Region: 50,000,000 | No | The total number of database objects that AWS DMS Fleet Advisor can process. A database object is any data structure used to store or reference data, including tables, views, stored procedures, functions, and triggers |
| Number of monitored objects in DMS Fleet Advisor collector | Each supported Region: 2,000 | No | Maximum number of monitored objects that a customer can add in a single DMS Fleet Advisor collector |
| Number of running Data Migrations | Each supported Region: 500 | Yes | Maximum number of running data migrations simultaneously. |
| Replication instances | Each supported Region: 60 | Yes | The maximum number of replication instances allowed in this account in the current Region. |
| Serverless replications | Each supported Region: 100 | Yes | The maximum number of serverless replications allowed in this account in the current Region. |
| Subnet groups | Each supported Region: 60 | Yes | The maximum number of subnet groups allowed in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Subnets per subnet group | Each supported Region: 60 | Yes | The maximum number of subnets allowed per subnet group. |
| Task count | Each supported Region: 600 | Yes | The maximum number of tasks allowed in this account in the current Region. |
| The amount of collected data in DMS Fleet Advisor | Each supported Region: 10 Gigabytes | No | The amount of data that can be collected by all DMS Fleet Advisor collectors |
| Total storage | Each supported Region: 30,000 Gigabytes | Yes | The maximum total storage (in GB) for all replication instances added together. |

# AWS Deadline Cloud endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | management.deadline.us-east-2.amazonaws.com<br>scheduling.deadline.us-east-2.amazonaws.com<br>scheduling.deadline.us-east-2.api.aws<br>management.deadline.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | management.deadline.us-east-1.amazonaws.com<br>scheduling.deadline.us-east-1.amazonaws.com<br>scheduling.deadline.us-east-1.api.aws<br>management.deadline.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | management.deadline.us-west-2.amazonaws.com<br>scheduling.deadline.us-west-2.amazonaws.com<br>scheduling.deadline.us-west-2.api.aws<br>management.deadline.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | management.deadline.ap-northeast-2.amazonaws.com<br>scheduling.deadline.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | scheduling.deadline.ap-northeast-2.api.aws | HTTPS |
| | | management.deadline.ap-northeast-2.api.aws | |
| Asia Pacific (Singapore) | ap-southeast-1 | management.deadline.ap-southeast-1.amazonaws.com | HTTPS |
| | | scheduling.deadline.ap-southeast-1.amazonaws.com | HTTPS |
| | | scheduling.deadline.ap-southeast-1.api.aws | HTTPS |
| | | management.deadline.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | management.deadline.ap-southeast-2.amazonaws.com | HTTPS |
| | | scheduling.deadline.ap-southeast-2.amazonaws.com | HTTPS |
| | | scheduling.deadline.ap-southeast-2.api.aws | HTTPS |
| | | management.deadline.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | management.deadline.ap-northeast-1.amazonaws.com | HTTPS |
| | | scheduling.deadline.ap-northeast-1.amazonaws.com | HTTPS |
| | | scheduling.deadline.ap-northeast-1.api.aws | HTTPS |
| | | management.deadline.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | management.deadline.eu-central-1.amazonaws.com<br>scheduling.deadline.eu-central-1.amazonaws.com<br>scheduling.deadline.eu-central-1.api.aws<br>management.deadline.eu-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | management.deadline.eu-west-1.amazonaws.com<br>scheduling.deadline.eu-west-1.amazonaws.com<br>scheduling.deadline.eu-west-1.api.aws<br>management.deadline.eu-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | management.deadline.eu-west-2.amazonaws.com<br>scheduling.deadline.eu-west-2.amazonaws.com<br>scheduling.deadline.eu-west-2.api.aws<br>management.deadline.eu-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Associated members per farm | Each supported Region: 75 | No | The maximum number of principals (users or groups) that can be associated to each farm in the current AWS Region. Each associated group counts as one member, regardless of how many users it contains. |
| Associated members per fleet | Each supported Region: 75 | No | The maximum number of principals (users or groups) that can be associated to each fleet in the current AWS Region. Each associated group counts as one member, regardless of how many users it contains. |
| Associated members per job | Each supported Region: 75 | No | The maximum number of principals (users or groups) that can be associated to each job in the current AWS Region. Each associated group counts as one member, regardless of how many users it contains. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Associated members per queue | Each supported Region: 75 | No | The maximum number of principals (users or groups) that can be associated to each queue in the current AWS Region. Each associated group counts as one member, regardless of how many users it contains. |
| Budgets per farm | Each supported Region: 20 | Yes | The maximum number of budgets per farm in the current AWS Region |
| Farms per region | Each supported Region: 2 | Yes | The maximum number of farms that can be created in the current AWS Region. |
| Fleets per farm | Each supported Region: 5 | Yes | The maximum number of fleets that can be created for each farm in the current AWS Region. |
| Jobs per farm | Each supported Region: 100,000 | Yes | The maximum number of jobs per farm in the current AWS Region. |
| License endpoints per region | Each supported Region: 5 | Yes | The maximum number of license endpoints in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| License sessions per license endpoint | Each supported Region: 500 | Yes | The maximum number of license sessions per license endpoint in the current AWS Region. |
| Limits per farm | Each supported Region: 50 | Yes | The maximum number of limits that can be created for each farm in the current AWS Region. |
| Monitors per region | Each supported Region: 1 | No | The maximum number of monitors in the current AWS Region. |
| OnDemand G instance GPUs per region | Each supported Region: 0 | Yes | The maximum number of on-demand G instance GPUs that can be provisioned across all service-managed fleets in the current AWS Region. |
| OnDemand vCPUs per region | Each supported Region: 50 | Yes | The maximum number of on-demand vCPUs that can be provisioned across all service-managed fleets in the current AWS Region. |
| Queue environments per queue | Each supported Region: 10 | No | The maximum number of queue environments that can be created for each queue in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Queue fleet associations per farm | Each supported Region: 100 | Yes | The maximum number of queue fleet associations per farm in the current AWS Region |
| Queue limit associations per queue | Each supported Region: 10 | Yes | The maximum number of limits that can be associated with each queue in the current AWS Region. |
| Queues per farm | Each supported Region: 20 | Yes | The maximum number of queues that can be created for each farm in the current AWS Region. |
| Resource configurations per fleet | Each supported Region: 1 | Yes | The maximum number of VPC Lattice resource configurations that can be added to each fleet. |
| Spot G Instance GPUs per region | Each supported Region: 0 | Yes | The maximum number of spot G instance GPUs that can be provisioned across all service-managed fleets in the current AWS Region. |
| Spot vCPUs per region | Each supported Region: 50 | Yes | The maximum number of spot vCPUs that can be provisioned across all service-managed fleets in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Step consumers per step | Each supported Region: 32 | Yes | The maximum number of steps that may declare a dependency on (consume) a single step within a job in the current AWS Region. |
| Step dependencies per step | Each supported Region: 128 | Yes | The maximum number of steps that a single step within a job may declare a dependency on in the current AWS Region. |
| Steps per job | Each supported Region: 200 | Yes | The maximum number of steps per job in the current AWS Region. |
| Storage for General Purpose SSD (gp3) volumes, in TiB | Each supported Region: 1 | Yes | The maximum aggregated amount of EBS storage, measured in TiB, that can be used across all fleets in the current AWS Region. |
| Storage for persistent General Purpose SSD (gp3) volumes, in TiB | Each supported Region: 4 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across persistent General Purpose SSD (gp3) volumes in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Storage profiles per farm | Each supported Region: 50 | No | The maximum number of storage profiles that can be created for each farm in the current AWS Region. |
| Tasks per chunk | Each supported Region: 150 | No | The maximum number of tasks that can be combined into a single chunk when submitting a job. |
| Tasks per job | Each supported Region: 10,000 | Yes | The maximum number of tasks per job in the current AWS Region. |
| Tasks per step | Each supported Region: 10,000 | Yes | The maximum number of tasks per step in the current AWS Region. |
| Wait-and-save vCPUs per region | Each supported Region: 50 | Yes | The maximum number of wait-and-save vCPUs that can be provisioned across all service-managed fleets in the current AWS Region. |
| Workers per farm | Each supported Region: 7,500 | Yes | The maximum number of workers per farm in the current AWS Region. |

For more information, see Quotas in the Deadline Cloud user guide.

# AWS DeepRacer endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Cars | Each supported Region: 20 | Yes | The maximum number of cars any account can have at same time. |
| Evaluation jobs | Each supported Region: 3 | Yes | The maximum number of concurrent evaluation jobs in same account in the same region. |
| Training jobs | Each supported Region: 4 | Yes | The maximum number of concurrent training jobs in same account in the same region. |

# Amazon Detective endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | api.detective.us-east-2.amazonaws.com | HTTPS |
| | | detective-fips.us-east-2.api.aws | HTTPS |
| | | api.detective-fips.us-east-2.amazonaws.com | HTTPS |
| | | detective.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | api.detective.us-east-1.amazonaws.com | HTTPS |
| | | detective-fips.us-east-1.api.aws | HTTPS |
| | | api.detective-fips.us-east-1.amazonaws.com | HTTPS |
| | | detective.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | api.detective.us-west-1.amazonaws.com | HTTPS |
| | | detective-fips.us-west-1.api.aws | HTTPS |
| | | api.detective-fips.us-west-1.amazonaws.com | HTTPS |
| | | detective.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | api.detective.us-west-2.amazonaws.com | HTTPS |
| | | detective-fips.us-west-2.api.aws | HTTPS |
| | | api.detective-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | detective.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | api.detective.af-south-1.amazonaws.com | HTTPS |
| | | detective.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | api.detective.ap-east-1.amazonaws.com | HTTPS |
| | | detective.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | api.detective.ap-south-1.amazonaws.com | HTTPS |
| | | detective.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | api.detective.ap-northeast-2.amazonaws.com | HTTPS |
| | | detective.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | api.detective.ap-southeast-1.amazonaws.com | HTTPS |
| | | detective.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | api.detective.ap-southeast-2.amazonaws.com | HTTPS |
| | | detective.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.detective.ap-northeast-1.amazonaws.com | HTTPS |
| | | detective.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | api.detective.ca-central-1.amazonaws.com | HTTPS |
| | | detective-fips.ca-central-1.api.aws | HTTPS |
| | | api.detective-fips.ca-central-1.amazonaws.com | HTTPS |
| | | detective.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | api.detective.eu-central-1.amazonaws.com | HTTPS |
| | | detective.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | api.detective.eu-west-1.amazonaws.com | HTTPS |
| | | detective.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | api.detective.eu-west-2.amazonaws.com | HTTPS |
| | | detective.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | api.detective.eu-south-1.amazonaws.com | HTTPS |
| | | detective.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | api.detective.eu-west-3.amazonaws.com | HTTPS |
| | | detective.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | api.detective.eu-north-1.amazonaws.com | HTTPS |
| | | detective.eu-north-1.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | api.detective.il-central-1.amazonaws.com | HTTPS |
| | | detective.il-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | api.detective.me-south-1.amazonaws.com<br>detective.me-south-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | api.detective.sa-east-1.amazonaws.com<br>detective.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | api.detective.us-gov-east-1.amazonaws.com<br>detective-fips.us-gov-east-1.api.aws<br>api.detective-fips.us-gov-east-1.amazonaws.com<br>detective.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | api.detective.us-gov-west-1.amazonaws.com<br>detective-fips.us-gov-west-1.api.aws<br>api.detective-fips.us-gov-west-1.amazonaws.com<br>detective.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

For more information, see Amazon Detective quotas.

# Amazon DevOps Guru endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | devops-guru.us-east-2.amazonaws.com | HTTPS |
| | | devops-guru-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | devops-guru.us-east-1.amazonaws.com | HTTPS |
| | | devops-guru-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | devops-guru.us-west-1.amazonaws.com | HTTPS |
| | | devops-guru-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | devops-guru.us-west-2.amazonaws.com | HTTPS |
| | | devops-guru-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | devops-guru.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | devops-guru.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | devops-guru.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | devops-guru.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | devops-guru.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | devops-guru.ca-central-1.amazonaws.com<br>devops-guru-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | devops-guru.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | devops-guru.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | devops-guru.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | devops-guru.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | devops-guru.eu-north-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | devops-guru.sa-east-1.amazonaws.com | HTTPS | |

### Service quotas

| Name | Default | Adjustable |
|---|---|---|
| Maximum number of Amazon SNS topics that you can specify at once | 2 | No |
| Maximum number of CloudFormation stacks that you can specify | 1,000 | No |
| Maximum number of Amazon SQS queues for monitoring | 100 | Yes |

For more information, see Quotas and limits in the *Amazon DevOps Guru User Guide*.

## AWS Device Farm endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | devicefarm.us-west-2.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrency for automation tests on metered devices | Each supported Region: 5 | Yes | The maximum number of concurrent metered devices running automation tests. |
| Concurrency for remote access on metered devices | Each supported Region: 2 | Yes | The maximum number of concurrent metered devices running remote access sessions. |
| Remote access session length in minutes | Each supported Region: 150 | No | The maximum length of a remote access session per device in minutes. |
| Test run timeout per device in minutes | Each supported Region: 150 | No | The maximum length of an automation test run per device in minutes. |
| Uploaded file size | Each supported Region: 4 Gigabytes | No | The maximum size of a file to be uploaded. |

# AWS Direct Connect endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | directconnect.us-east-2.amazonaws.com | HTTPS |
| | | directconnect-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | directconnect.us-east-1.amazonaws.com | HTTPS |
| | | directconnect-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | directconnect.us-west-1.amazonaws.com | HTTPS |
| | | directconnect-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | directconnect.us-west-2.amazonaws.com | HTTPS |
| | | directconnect-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | directconnect.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-east-1 | directconnect.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | directconnect.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | directconnect.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | directconnect.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | directconnect.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | directconnect.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | directconnect.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | directconnect.ap-northeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | directconnect.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | directconnect.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | directconnect.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | directconnect.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | directconnect.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | directconnect.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | directconnect.ca-central-1.amazonaws.com<br>directconnect-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | directconnect.ca-west-1.amazonaws.com<br>directconnect-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | directconnect.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | directconnect.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | directconnect.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | directconnect.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | directconnect.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | directconnect.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | directconnect.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | directconnect.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | directconnect.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | directconnect.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | directconnect.me-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | directconnect.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | directconnect.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | directconnect.us-gov-east-1.amazonaws.com<br>directconnect-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | directconnect.us-gov-west-1.amazonaws.com<br>directconnect-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Dedicated connections per location | Each supported Region: 10 | No | The maximum number of active AWS Direct Connect dedicated connections per location. |
| Direct Connect gateways per account | Each supported Region: 200 | No | The maximum number of AWS Direct Connect gateways per account. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Direct Connect gateways per transit gateway | Each supported Region: 20 | No | The maximum number of Direct Connect gateways associated to a transit gateway. |
| LAGs per Region | Each supported Region: 10 | No | The maximum number of link aggregation groups (LAGs) per AWS Region. |
| Members per LAG | Each supported Region: 4 | No | The maximum number of dedicated connections or interconnects per link aggregation group (LAG). |
| Prefixes per transit gateway | Each supported Region: 200 | No | The maximum number of prefixes per AWS Transit Gateway from AWS to on-premises on a transit virtual interface. |
| Private or public or transit virtual interfaces per hosted connection | Each supported Region: 1 | No | The maximum number of private, public, or transit virtual interfaces per AWS Direct Connect hosted connection. |
| Public or private virtual interfaces per dedicated connection | Each supported Region: 50 | No | The maximum number of private, or public interfaces per AWS Direct Connect dedicated connection. |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Routes per BGP session on private or transit virtual interfaces | Each supported Region: 100 | No | The maximum Routes per Border Gateway Protocol (BGP) on a private virtual interface or a transit virtual interface. |
| Routes per BGP session on public virtual interfaces | Each supported Region: 1,000 | No | The maximum routes per Border Gateway Protocol (BGP) on a public virtual interface |
| Total number of virtual interfaces per dedicated connection | Each supported Region: 54 | No | The maximum number of private, or public and transit virtual interfaces per AWS Direct Connect dedicated connection. |
| Transit gateways per AWS Direct Connect gateway | Each supported Region: 6 | No | The maximum number of transit gateways per AWS Direct Connect gateway. |
| Transit virtual interfaces per dedicated connection | Each supported Region: 4 | No | The maximum number of transit virtual interfaces per AWS Direct Connect dedicated connection. |
| Virtual interfaces per Direct Connect gateway | Each supported Region: 30 | No | The maximum number of virtual interfaces per AWS Direct Connect gateway. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Virtual interfaces per LAG | Each supported Region: 54 | No | The maximum number of virtual interfaces for each Link Aggregation Group (LAG) |
| Virtual private gateways per Direct Connect gateway | Each supported Region: 20 | No | The maximum number of virtual private gateways per AWS Direct Connect gateway. |

For more information, see Direct Connect Quotas in the *Direct Connect User Guide*.

# Directory Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Directory Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | ds.us-east-2.amazonaws.com | HTTPS | |
| | | ds-fips.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ds.us-east-1.amazonaws.com<br>ds-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | ds.us-west-1.amazonaws.com<br>ds-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | ds.us-west-2.amazonaws.com<br>ds-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | ds.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | ds.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | ds.ap-south-2.amazonaws.com | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | ds.ap-southeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | ds.ap-southeast-5.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | ds.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | ds.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ds.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ds.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ds.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | ds.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ds.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | ds.ap-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | ds.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ds.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | ds.ca-central-1.amazonaws.com<br>ds-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | ds.ca-west-1.amazonaws.com<br>ds-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | ds.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | ds.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | ds.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | ds.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | ds.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | ds.eu-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | ds.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | ds.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | ds.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | ds.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | ds.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | ds.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | ds.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | ds.us-gov-east-1.amazonaws.com<br>ds-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | ds.us-gov-west-1.amazonaws.com<br>ds-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

For a list of supported endpoints by directory type, see Region availability for AWS Directory Service.

## Directory Service Data endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ds-data.us-east-2.amazonaws.com<br>ds-data-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | ds-data.us-east-1.amazonaws.com<br>ds-data-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | ds-data.us-west-1.amazonaws.com<br>ds-data-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | ds-data.us-west-2.amazonaws.com<br>ds-data-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | ds-data.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ds-data.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | ds-data.ap-northeast-3.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | ds-data.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ds-data.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ds-data.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ds-data.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | ds-data.ca-central-1.amazonaws.com<br>ds-data-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | ds-data.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | ds-data.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | ds-data.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | ds-data.eu-west-3.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | ds-data.eu-north-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | ds-data.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AD Connector directories | Each supported Region: 10 | Yes | The maximum number of AD Connector directories that you can create in the current Region. |
| AWS Managed Microsoft AD directories | Each supported Region: 20 | Yes | The maximum number of AWS Managed Microsoft AD directories that you can create in the current region. |
| AWS Managed Microsoft AD domain controllers | Each supported Region: 20 | Yes | The maximum number of domain controllers that you can add to your AWS Managed Microsoft AD directory. |
| AWS Managed Microsoft AD manual snapshots | Each supported Region: 5 | No | The maximum number of manual snapshots that you can take for your |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | AWS Managed Microsoft AD directory. |
| Simple AD directories | Each supported Region: 10 | Yes | The maximum number of Simple AD directories that you can create in the current Region. |
| Simple AD manual snapshots | Each supported Region: 5 | No | The maximum number of manual snapshots that you can take for your Simple AD directory. |

For more information, see the following:

- AD Connector quotas
- AWS Managed Microsoft AD quotas
- Simple AD quotas

## Amazon DocumentDB endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | rds.us-east-2.amazonaws.com | HTTPS | |
| | | rds-fips.us-east-2.api.aws | HTTPS | |
| | | rds.us-east-2.api.aws | HTTPS | |
| | | rds-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | rds.us-east-1.amazonaws.com | HTTPS | |
| | | rds-fips.us-east-1.api.aws | HTTPS | |
| | | rds-fips.us-east-1.amazonaws.com | HTTPS | |
| | | rds.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | rds.us-west-2.amazonaws.com | HTTPS | |
| | | rds-fips.us-west-2.amazonaws.com | HTTPS | |
| | | rds.us-west-2.api.aws | HTTPS | |
| | | rds-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | rds.af-south-1.amazonaws.com | HTTPS | |
| | | rds.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | rds.ap-east-1.amazonaws.com | HTTPS | |
| | | rds.ap-east-1.api.aws | HTTPS | |
| Asia Pacific | ap-south-2 | rds.ap-south-2.amazonaws.com | HTTPS | |
| | | rds.ap-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rds.ap-southeast-3.amazonaws.com<br>rds.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rds.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rds.ap-southeast-4.amazonaws.com<br>rds.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rds.ap-south-1.amazonaws.com<br>rds.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | rds.ap-northeast-3.amazonaws.com<br>rds.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | rds.ap-northeast-2.amazonaws.com<br>rds.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | rds.ap-southeast-1.amazonaws.com<br>rds.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | rds.ap-southeast-2.amazonaws.com | HTTPS |
| | | rds.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | rds.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | rds.ap-northeast-1.amazonaws.com | HTTPS |
| | | rds.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | rds.ca-central-1.amazonaws.com | HTTPS |
| | | rds.ca-central-1.api.aws | HTTPS |
| | | rds-fips.ca-central-1.api.aws | HTTPS |
| | | rds-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | rds.ca-west-1.amazonaws.com | HTTPS |
| | | rds-fips.ca-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | rds.eu-central-1.amazonaws.com | HTTPS |
| | | rds.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | rds.eu-west-1.amazonaws.com | HTTPS |
| | | rds.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | rds.eu-west-2.amazonaws.com | HTTPS |
| | | rds.eu-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Milan) | eu-south-1 | rds.eu-south-1.amazonaws.com<br>rds.eu-south-1.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | rds.eu-west-3.amazonaws.com<br>rds.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | rds.eu-south-2.amazonaws.com<br>rds.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | rds.eu-north-1.amazonaws.com<br>rds.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | rds.eu-central-2.amazonaws.com<br>rds.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | rds.il-central-1.amazonaws.com<br>rds.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | rds.mx-central-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | rds.me-central-1.amazonaws.com<br>rds.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | rds.sa-east-1.amazonaws.com<br>rds.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | rds.us-gov-east-1.amazonaws.com<br>rds.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | rds.us-gov-west-1.amazonaws.com<br>rds.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

For information on finding and connecting to your cluster or instance endpoints, see Working with Amazon DocumentDB Endpoints in the *Amazon DocumentDB Developer Guide*.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Cluster parameter groups | Each supported Region: 50 | No | The maximum number of DB cluster parameter groups |
| Clusters | Each supported Region: 40 | Yes | The maximum number of clusters allowed in this account in the current Region |
| Event subscriptions | Each supported Region: 20 | Yes | The maximum number of event subscriptions |
| Instances | Each supported Region: 40 | Yes | The maximum number of DB instances allowed in this account in the current Region |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Manual cluster snapshots | Each supported Region: 100 | Yes | The maximum number of manual snapshots |
| Read replicas per cluster | Each supported Region: 15 | Yes | The maximum number of read replicas per cluster |
| Subnet groups | Each supported Region: 50 | Yes | The maximum number of DB subnet groups |
| Subnets per subnet group | Each supported Region: 20 | No | The maximum number of subnets per DB subnet group |
| Tags per resource | Each supported Region: 50 | No | The maximum number of tags per Amazon RDS resource |
| VPC security groups per instance | Each supported Region: 5 | No | The maximum number of DB security groups per Amazon VPC |

For more information, see Amazon DocumentDB Service Quotas in the *Amazon DocumentDB Developer Guide.*

## Amazon DynamoDB endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

For more information about this topic specific to DynamoDB, see Quotas in Amazon DynamoDB.

# Service endpoints

## DynamoDB

Newer versions of the AWS SDK connect to Amazon DynamoDB using the AWS-account-based endpoints listed below. For more information, see Account-based endpoints .

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | dynamodb.us-east-2.amazonaws.com | HTTP and HTTPS | |
| | | account-id.ddb.us-east-2.api.aws | HTTP and HTTPS | |
| | | dynamodb.us-east-2.api.aws | HTTP and HTTPS | |
| | | dynamodb-fips.us-east-2.api.aws | HTTPS | |
| | | dynamodb-fips.us-east-2.amazonaws.com | HTTPS | |
| | | account-id.ddb.us-east-2.amazonaws.com | HTTP and HTTPS | |
| US East (N. Virginia) | us-east-1 | dynamodb.us-east-1.amazonaws.com | HTTP and HTTPS | |
| | | dynamodb-fips.us-east-1.amazonaws.com | HTTPS | |
| | | account-id.ddb.us-east-1.api.aws | | |
| | | account-id.ddb.us-east-1.amazonaws.com | | |
| | | dynamodb.us-east-1.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | dynamodb-fips.us-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS<br><br>HTTP and HTTPS<br><br>HTTPS | |
| US West (N. California) | us-west-1 | dynamodb.us-west-1.amazonaws.com<br><br>dynamodb.us-west-1.api.aws<br><br>dynamodb-fips.us-west-1.api.aws<br><br>dynamodb-fips.us-west-1.amazonaws.com<br><br>account-id.ddb.us-west-1.amazonaws.com<br><br>account-id.ddb.us-west-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS<br><br>HTTPS<br><br>HTTPS<br><br>HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | dynamodb.us-west-2.amazonaws.com<br>account-id.ddb.us-west-2.amazonaws.com<br>dynamodb-fips.us-west-2.amazonaws.com<br>dynamodb.us-west-2.api.aws<br>account-id.ddb.us-west-2.api.aws<br>dynamodb-fips.us-west-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | dynamodb.af-south-1.amazonaws.com<br>account-id.ddb.af-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | dynamodb.ap-east-1.amazonaws.com<br>account-id.ddb.ap-east-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | dynamodb.ap-south-2.amazonaws.com<br>account-id.ddb.ap-south-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | dynamodb.ap-southeast-3.amazonaws.com<br>account-id.ddb.ap-southeast-3.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | dynamodb.ap-southeast-5.amazonaws.com<br>account-id.ddb.ap-southeast-5.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | dynamodb.ap-southeast-4.amazonaws.com<br>account-id.ddb.ap-southeast-4.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | dynamodb.ap-south-1.amazonaws.com<br>account-id.ddb.ap-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | dynamodb.ap-southeast-6.amazonaws.com<br>account-id.ddb.ap-southeast-6.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | dynamodb.ap-northeast-3.amazonaws.com<br>account-id.ddb.ap-northeast-3.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | dynamodb.ap-northeast-2.amazonaws.com<br>account-id.ddb.ap-northeast-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | dynamodb.ap-southeast-1.amazonaws.com<br>account-id.ddb.ap-southeast-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | dynamodb.ap-southeast-2.amazonaws.com<br>account-id.ddb.ap-southeast-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | dynamodb.ap-east-2.amazonaws.com<br>account-id.ddb.ap-east-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | dynamodb.ap-southeast-7.amazonaws.com<br>account-id.ddb.ap-southeast-7.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | dynamodb.ap-northeast-1.amazonaws.com<br>account-id.ddb.ap-northeast-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | dynamodb.ca-central-1.amazonaws.com<br>account-id.ddb.ca-central-1.amazonaws.com<br>dynamodb-fips.ca-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | dynamodb.ca-west-1.amazonaws.com<br>dynamodb-fips.ca-west-1.amazonaws.com<br>account-id.ddb.ca-west-1.amazonaws.com | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | dynamodb.eu-central-1.amazonaws.com<br>account-id.ddb.eu-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | dynamodb.eu-west-1.amazonaws.com<br>account-id.ddb.eu-west-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (London) | eu-west-2 | dynamodb.eu-west-2.amazonaws.com<br>account-id.ddb.eu-west-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Milan) | eu-south-1 | dynamodb.eu-south-1.amazonaws.com<br>account-id.ddb.eu-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | dynamodb.eu-west-3.amazonaws.com<br>account-id.ddb.eu-west-3.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | dynamodb.eu-south-2.amazonaws.com<br>account-id.ddb.eu-south-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | dynamodb.eu-north-1.amazonaws.com<br>account-id.ddb.eu-north-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Zurich) | eu-central-2 | dynamodb.eu-central-2.amazonaws.com<br>account-id.ddb.eu-central-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Israel (Tel Aviv) | il-central-1 | dynamodb.il-central-1.amazonaws.com<br>account-id.ddb.il-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | dynamodb.mx-central-1.amazonaws.com<br>account-id.ddb.mx-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | dynamodb.me-south-1.amazonaws.com<br>account-id.ddb.me-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| Middle East (UAE) | me-central-1 | dynamodb.me-central-1.amazonaws.com<br>account-id.ddb.me-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | dynamodb.sa-east-1.amazonaws.com<br>account-id.ddb.sa-east-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | dynamodb.us-gov-east-1.amazonaws.com<br>dynamodb-fips.us-gov-east-1.amazonaws.com<br>dynamodb.us-gov-east-1.api.aws<br>dynamodb-fips.us-gov-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | dynamodb.us-gov-west-1.amazonaws.com<br>dynamodb-fips.us-gov-west-1.api.aws<br>dynamodb.us-gov-west-1.api.aws<br>dynamodb-fips.us-gov-west-1.amazonaws.com | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS<br>HTTPS | |

## DynamoDB Accelerator (DAX)

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | dax.us-east-2.amazonaws.com<br>dax.us-east-2.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | dax.us-east-1.amazonaws.com<br>dax.us-east-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | dax.us-west-1.amazonaws.com<br>dax.us-west-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | dax.us-west-2.amazonaws.com<br>dax.us-west-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | dax.ap-south-2.amazonaws.com<br>dax.ap-south-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | dax.ap-south-1.amazonaws.com<br>dax.ap-south-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | dax.ap-southeast-1.amazonaws.com<br>dax.ap-southeast-1.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | dax.ap-southeast-2.amazonaws.com<br>dax.ap-southeast-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | dax.ap-northeast-1.amazonaws.com<br>dax.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | dax.eu-central-1.amazonaws.com<br>dax.eu-central-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | dax.eu-west-1.amazonaws.com<br>dax.eu-west-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | dax.eu-west-2.amazonaws.com<br>dax.eu-west-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | dax.eu-west-3.amazonaws.com<br>dax.eu-west-3.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | dax.eu-south-2.amazonaws.com<br>dax.eu-south-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | dax.eu-north-1.amazonaws.com<br>dax.eu-north-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | dax.sa-east-1.amazonaws.com<br>dax.sa-east-1.api.aws | HTTP and HTTPS<br>HTTPS | |

## Amazon DynamoDB Streams

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | streams.dynamodb.us-east-2.amazonaws.com<br>streams-dynamodb-fips.us-east-2.api.aws<br>streams.dynamodb-fips.us-east-2.amazonaws.com<br>streams-dynamodb.us-east-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | streams.dynamodb.us-east-1.amazonaws.com<br>streams-dynamodb-fips.us-east-1.api.aws<br>streams.dynamodb-fips.us-east-1.amazonaws.com<br>streams-dynamodb.us-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (N. California) | us-west-1 | streams.dynamodb.us-west-1.amazonaws.com<br>streams-dynamodb-fips.us-west-1.api.aws<br>streams.dynamodb-fips.us-west-1.amazonaws.com<br>streams-dynamodb.us-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | streams.dynamodb.us-west-2.amazonaws.com<br>streams-dynamodb-fips.us-west-2.api.aws<br>streams.dynamodb-fips.us-west-2.amazonaws.com<br>streams-dynamodb.us-west-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | streams.dynamodb.af-south-1.amazonaws.com<br>streams-dynamodb.af-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | streams.dynamodb.ap-east-1.amazonaws.com<br>streams-dynamodb.ap-east-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | streams.dynamodb.ap-south-2.amazonaws.com<br>streams-dynamodb.ap-south-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | streams.dynamodb.ap-southeast-3.amazonaws.com<br>streams-dynamodb.ap-southeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | streams.dynamodb.ap-southeast-5.amazonaws.com<br><br>streams-dynamodb.ap-southeast-5.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | streams.dynamodb.ap-southeast-4.amazonaws.com<br><br>streams-dynamodb.ap-southeast-4.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | streams.dynamodb.ap-south-1.amazonaws.com<br><br>streams-dynamodb.ap-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | streams.dynamodb.ap-southeast-6.amazonaws.com<br><br>streams-dynamodb.ap-southeast-6.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | streams.dynamodb.ap-northeast-3.amazonaws.com<br><br>streams-dynamodb.ap-northeast-3.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | streams.dynamodb.ap-northeast-2.amazonaws.com<br><br>streams-dynamodb.ap-northeast-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | streams.dynamodb.ap-southeast-1.amazonaws.com<br><br>streams-dynamodb.ap-southeast-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | streams.dynamodb.ap-southeast-2.amazonaws.com<br><br>streams-dynamodb.ap-southeast-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | streams.dynamodb.ap-east-2.amazonaws.com<br>streams-dynamodb.ap-east-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | streams.dynamodb.ap-southeast-7.amazonaws.com<br>streams-dynamodb.ap-southeast-7.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | streams.dynamodb.ap-northeast-1.amazonaws.com<br>streams-dynamodb.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | streams.dynamodb.ca-central-1.amazonaws.com<br>streams-dynamodb-fips.ca-central-1.api.aws<br>streams.dynamodb-fips.ca-central-1.amazonaws.com<br>streams-dynamodb.ca-central-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | streams.dynamodb.ca-west-1.amazonaws.com<br>streams-dynamodb-fips.ca-west-1.api.aws<br>streams.dynamodb-fips.ca-west-1.amazonaws.com<br>streams-dynamodb.ca-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | streams.dynamodb.eu-central-1.amazonaws.com<br>streams-dynamodb.eu-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | streams.dynamodb.eu-west-1.amazonaws.com<br>streams-dynamodb.eu-west-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (London) | eu-west-2 | streams.dynamodb.eu-west-2.amazonaws.com<br>streams-dynamodb.eu-west-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | streams.dynamodb.eu-south-1.amazonaws.com<br>streams-dynamodb.eu-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | streams.dynamodb.eu-west-3.amazonaws.com<br>streams-dynamodb.eu-west-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | streams.dynamodb.eu-south-2.amazonaws.com<br>streams-dynamodb.eu-south-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | streams.dynamodb.eu-north-1.amazonaws.com<br>streams-dynamodb.eu-north-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | streams.dynamodb.eu-central-2.amazonaws.com<br><br>streams-dynamodb.eu-central-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Israel (Tel Aviv) | il-central-1 | streams.dynamodb.il-central-1.amazonaws.com<br><br>streams-dynamodb.il-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Mexico (Central) | mx-central-1 | streams.dynamodb.mx-central-1.amazonaws.com<br><br>streams-dynamodb.mx-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | streams.dynamodb.me-south-1.amazonaws.com<br><br>streams-dynamodb.me-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | streams.dynamodb.me-central-1.amazonaws.com<br>streams-dynamodb.me-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | streams.dynamodb.sa-east-1.amazonaws.com<br>streams-dynamodb.sa-east-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | streams.dynamodb.us-gov-east-1.amazonaws.com<br>streams-dynamodb-fips.us-gov-east-1.api.aws<br>streams.dynamodb-fips.us-gov-east-1.amazonaws.com<br>streams-dynamodb.us-gov-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | streams.dynamodb.us-gov-west-1.amazonaws.com<br><br>streams-dynamodb-fips.us-gov-west-1.api.aws<br><br>streams.dynamodb-fips.us-gov-west-1.amazonaws.com<br><br>streams-dynamodb.us-gov-west-1.api.aws | HTTP and HTTPS<br><br>HTTPS<br><br>HTTPS<br><br>HTTP and HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Account-level read throughput limit (Provisioned mode) | Each supported Region: 80,000 | Yes | The maximum number of read capacity units allocated for the account; applicable only for tables (including all associated global secondary indexes) in provisioned read/write capacity mode. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#default-limits-throughput-capacity-modes |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Account-level write throughput limit (Provisioned mode) | Each supported Region: 80,000 | Yes | The maximum number of write capacity units allocated for the account; applicable only for tables (including all associated global secondary indexes) in provisioned read/write capacity mode. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#default-limits-throughput-capacity-modes |
| Concurrent control plane operations | Each supported Region: 500 | Yes | The maximum number of allowed concurrent control plane operations. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#limits-api |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Global Secondary Indexes per table | Each supported Region: 20 | Yes | The maximum number of global secondary indexes that can be created for a table. For more information, see, https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-secondary-indexes |
| Maximum Incremental Export concurrent data size | Each supported Region: 100 Terabytes | Yes | Maximum limit on size of in flight incremental export requests. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-table-export |
| Maximum Incremental Export concurrent requests | Each supported Region: 300 | Yes | Maximum limit on number of in flight incremental export requests. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-table-export |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum Incremental Export period window | Each supported Region: 24 | No | The limit on the maximum export period (in hours) for an incremental export request. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-table-export |
| Maximum number of tables | Each supported Region: 2,500 | Yes | The maximum number of tables that can be created per region. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-tables |
| Minimum Incremental Export period window | Each supported Region: 15 | No | The limit on the minimum export period (in minutes) for an incremental export request. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-table-export |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Provisioned capacity decreases per day | Each supported Region: 27 | Yes | A decrease is allowed up to four times any time per day (GMT time zone). Also, if there was no decrease in the past hour, an additional decrease is allowed, effectively bringing the maximum number of decreases in a day to 27 times. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html |
| Table-level read throughput limit | Each supported Region: 40,000 | Yes | The maximum number of read throughput allocated for a table or global secondary index. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#default-limits-throughput-capacity-modes |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Table-level write throughput limit | Each supported Region: 40,000 | Yes | The maximum number of write throughput allocated for a table or global secondary index. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html#default-limits-throughput-capacity-modes |
| Write throughput limit for DynamoDB Streams (Provisioned mode) | Each supported Region: 40,000 | Yes | The maximum number of write capacity units allowed for a table with streams enabled; applicable only for tables in provisioned read/write capacity mode. Other quotas might also apply. For more information, see https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-dynamodb-streams |

DAX has the following quotas.

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Nodes per cluster | Each supported Region: 11 | No | The maximum number of nodes per cluster, including the primary node as well as any read replica nodes. |
| Parameter groups | Each supported Region: 20 | No | The maximum number of parameter groups in a single AWS region. |
| Subnet groups | Each supported Region: 50 | No | The maximum number of subnet groups in a single AWS region. |
| Subnets per subnet group | Each supported Region: 20 | No | The maximum number of subnets per subnet group in a single AWS region. |
| Total number of nodes | Each supported Region: 50 | Yes | The maximum total number of nodes per AWS account in a single AWS region. |

# AWS Elastic Beanstalk endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

> **Note**
>
> The AWS Elastic Beanstalk service offers dual stack endpoints, so that you can send it either IPv4 or IPv6 requests. For more information about the naming syntax for dual-stack endpoints, see Dual stack endpoints in this guide. For more information about Elastic Beanstalk support for dual-stack endpoints, see IPv6 support in the *AWS Elastic Beanstalk Developer Guide*.

## Elastic Beanstalk

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | elasticbeanstalk.us-east-2.amazonaws.com<br>elasticbeanstalk-fips.us-east-2.api.aws<br>elasticbeanstalk-fips.us-east-2.amazonaws.com<br>elasticbeanstalk.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | Z14LCN19Q5QHIC |
| US East (N. Virginia) | us-east-1 | elasticbeanstalk.us-east-1.amazonaws.com<br>elasticbeanstalk-fips.us-east-1.api.aws<br>elasticbeanstalk-fips.us-east-1.amazonaws.com<br>elasticbeanstalk.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | Z117KPS5GTRQ2G |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| US West (N. California) | us-west-1 | elasticbeanstalk.us-west-1.amazonaws.com<br>elasticbeanstalk-fips.us-west-1.api.aws<br>elasticbeanstalk-fips.us-west-1.amazonaws.com<br>elasticbeanstalk.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | Z1LQECGX5PH1X |
| US West (Oregon) | us-west-2 | elasticbeanstalk.us-west-2.amazonaws.com<br>elasticbeanstalk-fips.us-west-2.api.aws<br>elasticbeanstalk-fips.us-west-2.amazonaws.com<br>elasticbeanstalk.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | Z38NKT9BP95V3O |
| Africa (Cape Town) | af-south-1 | elasticbeanstalk.af-south-1.amazonaws.com<br>elasticbeanstalk.af-south-1.api.aws | HTTPS<br>HTTPS | Z1EI3BVKMKK4AM |
| Asia Pacific (Hong Kong) | ap-east-1 | elasticbeanstalk.ap-east-1.amazonaws.com<br>elasticbeanstalk.ap-east-1.api.aws | HTTPS<br>HTTPS | ZPWYUBWRU171A |
| Asia Pacific (Hyderabad) | ap-south-2 | elasticbeanstalk.ap-south-2.amazonaws.com<br>elasticbeanstalk.ap-south-2.api.aws | HTTPS<br>HTTPS | Z10223522IBWPBF2C2FJS |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | elasticbeanstalk.ap-southeast-3.amazonaws.com<br>elasticbeanstalk.ap-southeast-3.api.aws | HTTPS<br>HTTPS | Z05913172VM7EAZB40TA8 |
| Asia Pacific (Malaysia) | ap-southeast-5 | elasticbeanstalk.ap-southeast-5.amazonaws.com<br>elasticbeanstalk.ap-southeast-5.api.aws | HTTPS<br>HTTPS | Z01812971H0QCYWSL7WOH |
| Asia Pacific (Melbourne) | ap-southeast-4 | elasticbeanstalk.ap-southeast-4.amazonaws.com<br>elasticbeanstalk.ap-southeast-4.api.aws | HTTPS<br>HTTPS | Z0666869LC74UHAO5YE4 |
| Asia Pacific (Mumbai) | ap-south-1 | elasticbeanstalk.ap-south-1.amazonaws.com<br>elasticbeanstalk.ap-south-1.api.aws | HTTPS<br>HTTPS | Z18NTBI3Y7N9TZ |
| Asia Pacific (New Zealand) | ap-southeast-6 | elasticbeanstalk.ap-southeast-6.amazonaws.com<br>elasticbeanstalk.ap-southeast-6.api.aws | HTTPS<br>HTTPS | Z01144401H1NECCLJDD4D |
| Asia Pacific (Osaka) | ap-northeast-3 | elasticbeanstalk.ap-northeast-3.amazonaws.com<br>elasticbeanstalk.ap-northeast-3.api.aws | HTTPS<br>HTTPS | ZNE5GEY1TIAGY |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | elasticbeanstalk.ap-northeast-2.amazonaws.com<br>elasticbeanstalk.ap-northeast-2.api.aws | HTTPS<br>HTTPS | Z3JE5OI70TWKCP |
| Asia Pacific (Singapore) | ap-southeast-1 | elasticbeanstalk.ap-southeast-1.amazonaws.com<br>elasticbeanstalk.ap-southeast-1.api.aws | HTTPS<br>HTTPS | Z16FZ9L249IFLT |
| Asia Pacific (Sydney) | ap-southeast-2 | elasticbeanstalk.ap-southeast-2.amazonaws.com<br>elasticbeanstalk.ap-southeast-2.api.aws | HTTPS<br>HTTPS | Z2PCDNR3VC2G1N |
| Asia Pacific (Thailand) | ap-southeast-7 | elasticbeanstalk.ap-southeast-7.amazonaws.com<br>elasticbeanstalk.ap-southeast-7.api.aws | HTTPS<br>HTTPS | Z08384933QM5LSQCVMNZM |
| Asia Pacific (Tokyo) | ap-northeast-1 | elasticbeanstalk.ap-northeast-1.amazonaws.com<br>elasticbeanstalk.ap-northeast-1.api.aws | HTTPS<br>HTTPS | Z1R25G3KIG2GBW |
| Canada (Central) | ca-central-1 | elasticbeanstalk.ca-central-1.amazonaws.com<br>elasticbeanstalk.ca-central-1.api.aws | HTTPS<br>HTTPS | ZJFCZL7SSZB5I |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | elasticbeanstalk.ca-west-1.amazonaws.com<br>elasticbeanstalk.ca-west-1.api.aws | HTTPS<br>HTTPS | Z1021028356Y0CYS11DQI |
| Europe (Frankfurt) | eu-central-1 | elasticbeanstalk.eu-central-1.amazonaws.com<br>elasticbeanstalk.eu-central-1.api.aws | HTTPS<br>HTTPS | Z1FRNW7UH4DEZJ |
| Europe (Ireland) | eu-west-1 | elasticbeanstalk.eu-west-1.amazonaws.com<br>elasticbeanstalk.eu-west-1.api.aws | HTTPS<br>HTTPS | Z2NYPWQ7DFZAZH |
| Europe (London) | eu-west-2 | elasticbeanstalk.eu-west-2.amazonaws.com<br>elasticbeanstalk.eu-west-2.api.aws | HTTPS<br>HTTPS | Z1GKAAAUGATPF1 |
| Europe (Milan) | eu-south-1 | elasticbeanstalk.eu-south-1.amazonaws.com<br>elasticbeanstalk.eu-south-1.api.aws | HTTPS<br>HTTPS | Z10VDYYOA2JFKM |
| Europe (Paris) | eu-west-3 | elasticbeanstalk.eu-west-3.amazonaws.com<br>elasticbeanstalk.eu-west-3.api.aws | HTTPS<br>HTTPS | Z5WN6GAYWG5OB |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | elasticbeanstalk.eu-south-2.amazonaws.com<br>elasticbeanstalk.eu-south-2.api.aws | HTTPS<br>HTTPS | Z0243492AO4B9S3KI68O |
| Europe (Stockholm) | eu-north-1 | elasticbeanstalk.eu-north-1.amazonaws.com<br>elasticbeanstalk.eu-north-1.api.aws | HTTPS<br>HTTPS | Z23GO28BZ5AETM |
| Europe (Zurich) | eu-central-2 | elasticbeanstalk.eu-central-2.amazonaws.com<br>elasticbeanstalk.eu-central-2.api.aws | HTTPS<br>HTTPS | Z00227012FSHBMZNNSJJI |
| Israel (Tel Aviv) | il-central-1 | elasticbeanstalk.il-central-1.amazonaws.com<br>elasticbeanstalk.il-central-1.api.aws | HTTPS<br>HTTPS | Z02941091PERNCB1MI5H7 |
| Middle East (Bahrain) | me-south-1 | elasticbeanstalk.me-south-1.amazonaws.com<br>elasticbeanstalk.me-south-1.api.aws | HTTPS<br>HTTPS | Z2BBTEKR2I36N2 |
| Middle East (UAE) | me-central-1 | elasticbeanstalk.me-central-1.amazonaws.com<br>elasticbeanstalk.me-central-1.api.aws | HTTPS<br>HTTPS | Z0650511N84AAKGW7QUK |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | elasticbeanstalk.sa-east-1.amazonaws.com<br>elasticbeanstalk.sa-east-1.api.aws | HTTPS<br>HTTPS | Z10X7K2B4QSOFV |
| AWS GovCloud (US-East) | us-gov-east-1 | elasticbeanstalk.us-gov-east-1.amazonaws.com<br>elasticbeanstalk.us-gov-east-1.api.aws | HTTPS<br>HTTPS | Z35TSARG0EJ4VU |
| AWS GovCloud (US-West) | us-gov-west-1 | elasticbeanstalk.us-gov-west-1.amazonaws.com<br>elasticbeanstalk.us-gov-west-1.api.aws | HTTPS<br>HTTPS | Z4KAURWC4UUUG |

## Elastic Beanstalk Health Service

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | elasticbeanstalk-health.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | elasticbeanstalk-health.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | elasticbeanstalk-health.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | elasticbeanstalk-health.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | elasticbeanstalk-health.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | elasticbeanstalk-health.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | elasticbeanstalk-health.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | elasticbeanstalk-health.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | elasticbeanstalk-health.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | elasticbeanstalk-health.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | elasticbeanstalk-health.ap-southeast-6.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | elasticbeanstalk-health.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | elasticbeanstalk-health.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | elasticbeanstalk-health.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | elasticbeanstalk-health.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | elasticbeanstalk-health.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | elasticbeanstalk-health.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | elasticbeanstalk-health.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | elasticbeanstalk-health.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | elasticbeanstalk-health.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | elasticbeanstalk-health.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | elasticbeanstalk-health.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | elasticbeanstalk-health.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | elasticbeanstalk-health.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | elasticbeanstalk-health.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | elasticbeanstalk-health.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | elasticbeanstalk-health.il-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | elasticbeanstalk-health.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | elasticbeanstalk-health.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | elasticbeanstalk-health.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | elasticbeanstalk-health.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | elasticbeanstalk-health.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Application versions | Each supported Region: 1,000 | Yes | The maximum number of application versions that you can create in this account in the current Region. The limit applies across applications, not per application. |
| Applications | Each supported Region: 75 | Yes | The maximum number of applications that you can create in this account in the current Region. |
| Configuration templates | Each supported Region: 2,000 | Yes | The maximum number of configuration templates that you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Custom platform versions | Each supported Region: 50 | Yes | The maximum number of custom platform versions that you can create in this account in the current Region. The limit applies across custom platforms, not per custom platform. |
| Environments | Each supported Region: 200 | Yes | The maximum number of environments that you can create in this account in the current Region. The limit applies across applications, not per application. |

# Amazon Elastic Block Store endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Topics

- Endpoints for Amazon EBS in Amazon EC2
- Endpoints for the EBS direct APIs

## Endpoints for Amazon EBS in Amazon EC2

The API actions to manage EBS volumes, snapshots, and encryption are part of the Amazon EC2 API. For more information, see Amazon EBS actions in the *Amazon EC2 API Reference*.

For the service endpoints for Amazon EC2, see the section called "Amazon EC2".

## Endpoints for the EBS direct APIs

Use the EBS direct APIs endpoints to directly read the data on your Amazon EBS snapshots, and identify the difference between two snapshots. For more information, see Use EBS direct APIs to access the contents of an Amazon EBS snapshot in the *Amazon Elastic Compute Cloud User Guide*.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ebs.us-east-2.amazonaws.com | HTTPS |
| | | ebs-fips.us-east-2.api.aws | HTTPS |
| | | ebs.us-east-2.api.aws | HTTPS |
| | | ebs-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | ebs.us-east-1.amazonaws.com | HTTPS |
| | | ebs-fips.us-east-1.api.aws | HTTPS |
| | | ebs-fips.us-east-1.amazonaws.com | HTTPS |
| | | ebs.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ebs.us-west-1.amazonaws.com | HTTPS |
| | | ebs.us-west-1.api.aws | HTTPS |
| | | ebs-fips.us-west-1.api.aws | HTTPS |
| | | ebs-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | ebs.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | ebs-fips.us-west-2.api.aws | HTTPS |
| | | ebs.us-west-2.api.aws | HTTPS |
| | | ebs-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | ebs.af-south-1.amazonaws.com | HTTPS |
| | | ebs.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | ebs.ap-east-1.amazonaws.com | HTTPS |
| | | ebs.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | ebs.ap-south-2.amazonaws.com | HTTPS |
| | | ebs.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | ebs.ap-southeast-3.amazonaws.com | HTTPS |
| | | ebs.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | ebs.ap-southeast-5.amazonaws.com | HTTPS |
| | | ebs.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | ebs.ap-southeast-4.amazonaws.com | HTTPS |
| | | ebs.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | ebs.ap-south-1.amazonaws.com<br>ebs.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ebs.ap-southeast-6.amazonaws.com<br>ebs.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ebs.ap-northeast-3.amazonaws.com<br>ebs.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ebs.ap-northeast-2.amazonaws.com<br>ebs.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | ebs.ap-southeast-1.amazonaws.com<br>ebs.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ebs.ap-southeast-2.amazonaws.com<br>ebs.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | ebs.ap-east-2.amazonaws.com<br>ebs.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | ebs.ap-southeast-7.amazonaws.com<br>ebs.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | ebs.ap-northeast-1.amazonaws.com | HTTPS |
| | | ebs.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | ebs.ca-central-1.amazonaws.com | HTTPS |
| | | ebs.ca-central-1.api.aws | HTTPS |
| | | ebs-fips.ca-central-1.amazonaws.com | HTTPS |
| | | ebs-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | ebs.ca-west-1.amazonaws.com | HTTPS |
| | | ebs.ca-west-1.api.aws | HTTPS |
| | | ebs-fips.ca-west-1.amazonaws.com | HTTPS |
| | | ebs-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ebs.eu-central-1.amazonaws.com | HTTPS |
| | | ebs.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ebs.eu-west-1.amazonaws.com | HTTPS |
| | | ebs.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | ebs.eu-west-2.amazonaws.com | HTTPS |
| | | ebs.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | ebs.eu-south-1.amazonaws.com | HTTPS |
| | | ebs.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | ebs.eu-west-3.amazonaws.com | HTTPS |
| | | ebs.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | ebs.eu-south-2.amazonaws.com | HTTPS | |
| | | ebs.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | ebs.eu-north-1.amazonaws.com | HTTPS | |
| | | ebs.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | ebs.eu-central-2.amazonaws.com | HTTPS | |
| | | ebs.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | ebs.il-central-1.amazonaws.com | HTTPS | |
| | | ebs.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | ebs.mx-central-1.amazonaws.com | HTTPS | |
| | | ebs.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | ebs.me-south-1.amazonaws.com | HTTPS | |
| | | ebs.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | ebs.me-central-1.amazonaws.com | HTTPS | |
| | | ebs.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | ebs.sa-east-1.amazonaws.com | HTTPS | |
| | | ebs.sa-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | ebs.us-gov-east-1.amazonaws.com | HTTPS | |
| | | ebs.us-gov-east-1.api.aws | HTTPS | |
| | | ebs-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | ebs-fips.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | ebs.us-gov-west-1.amazonaws.com | HTTPS | |
| | | ebs-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | ebs-fips.us-gov-west-1.api.aws | HTTPS | |
| | | ebs.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

Your quotas can change over time. Amazon EBS constantly monitors your provisioned storage and IOPS usage within each Region and might automatically increase your quotas, on a per-Region basis, based on your usage.

Even though Amazon EBS automatically increases your quotas based on your usage, you can request a quota increase if needed. For example, if you intend to use more gp3 storage in US East (N. Virginia) than your current quota allows, you can request a quota increase for that volume type in that Region. To request a quota increase, use the AWS Service Quotas console.

### Considerations

- The quota for **Concurrent snapshot copies per destination Region** is not adjustable using Service Quotas. However, you can request an increase for this quota by contacting AWS Support.
- The **IOPS modifications and Storage modifications** quotas apply to the aggregated current value (for *size* or *IOPS*, depending on the quota) of volumes that can undergo modifications concurrently. You can make concurrent modification requests for volumes that have combined current value (for *size* or *IOPS*) up to the quota. For example, if your **IOPS modifications for Provisioned IOPS SSD (io1) volumes** quota is `50,000`, you can make concurrent IOPS

modifications requests for any number of `io1` volumes as long as their combined current IOPS is equal to or less than `50,000`. If you have three `io1` volumes provisioned with `20,000` IOPS each, you can request IOPS modifications for two volumes concurrently (`20,000 * 2 < 50,000`). If you submit a concurrent IOPS modification request for the third volume, you exceed your quota and that request fails (`20,000 * 3 > 50,000`).

- Amazon EBS has the following non-adjustable limits for the number of EBS volumes per instance launch request.
  - `2500` — `us-east-1`, `us-west-2`, `eu-west-1`, and `ap-northeast-1`
  - `500` — all other Regions

  This limit applies to instance launch requests that you make, and to instance launch requests made by AWS services, such as Amazon EMR, on your behalf. If your instance launch request fails as a result of exceeding this limit, we recommend that you adjust the EBS volume configuration in the launch request to ensure the number of volumes is below the limit, or that you work with your technical account manager (TAM) to explore other options for launching your cluster without exceeding the limit.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Archived snapshots per volume | Each supported Region: 25 | Yes | The maximum number of archived snapshots per volume. |
| CompleteSnapshot requests per account | Each supported Region: 10 per second | No | The maximum number of CompleteSnapshot requests allowed per account. |
| Concurrent snapshot copies per destination Region | Each supported Region: 20 | No | The maximum number of concurrent snapshot copies to a single destination Region. |
| Concurrent snapshots per Cold HDD (sc1) volume | Each supported Region: 1 | No | The maximum number of concurrent snapshots per |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | Cold HDD (sc1) volume in this Region. |
| Concurrent snapshots per General Purpose SSD (gp2) volume | Each supported Region: 5 | No | The maximum number of concurrent snapshots per General Purpose SSD (gp2) volume in this Region. |
| Concurrent snapshots per General Purpose SSD (gp3) volume | Each supported Region: 5 | No | The maximum number of concurrent snapshots per General Purpose SSD (gp3) volume in this Region. |
| Concurrent snapshots per Magnetic (standard) volume | Each supported Region: 5 | No | The maximum number of concurrent snapshots per Magnetic (standard) volume in this Region. |
| Concurrent snapshots per Provisioned IOPS SSD (io1) volume | Each supported Region: 5 | No | The maximum number of concurrent snapshots per Provisioned IOPS SSD (io1) volume in this Region. |
| Concurrent snapshots per Provisioned IOPS SSD (io2) volume | Each supported Region: 5 | No | The maximum number of concurrent snapshots per Provisioned IOPS SSD (io2) volume in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrent snapshots per Throughput Optimized HDD (st1) volume | Each supported Region: 1 | No | The maximum number of concurrent snapshots per Throughput Optimized HDD (st1) volume in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Fast snapshot restore | us-east-1: 5<br>us-east-2: 5<br>us-west-1: 5<br>us-west-2: 5<br>af-south-1: 5<br>ap-east-1: 5<br>ap-east-2: 5<br>ap-northeast-1: 5<br>ap-northeast-2: 5<br>ap-northeast-3: 5<br>ap-south-1: 5<br>ap-south-2: 5<br>ap-southeast-1: 5<br>ap-southeast-2: 5<br>ap-southeast-3: 5<br>ap-southeast-4: 5<br>ap-southeast-5: 5<br>ap-southeast-6: 5<br>ap-southeast-7: 5<br>ca-central-1: 5 | Yes | The maximum number of snapshots that can be enabled for fast snapshot restore in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | ca-west-1: 5<br>eu-central-1: 5<br>eu-central-2: 5<br>eu-north-1: 5<br>eu-south-1: 5<br>eu-south-2: 5<br>eu-west-1: 5<br>eu-west-2: 5<br>eu-west-3: 5<br>il-central-1: 5<br>mx-central-1: 5<br>sa-east-1: 5<br>Each of the other supported Regions: 5 | | |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetSnapshotBlock requests per account | us-east-1: 5,000 per second<br><br>us-east-2: 5,000 per second<br><br>us-west-2: 5,000 per second<br><br>ap-southeast-1: 5,000 per second<br><br>eu-west-1: 5,000 per second<br><br>Each of the other supported Regions: 1,000 per second | Yes | The maximum number of GetSnapshotBlock requests allowed per account. |
| GetSnapshotBlock requests per snapshot | Each supported Region: 1,000 per second | No | The maximum number of GetSnapshotBlock requests allowed per snapshot. |
| IOPS for Provisioned IOPS SSD (io1) volumes | Each supported Region: 300,000 | Yes | The maximum aggregated number of IOPS that can be provisioned across Provisioned IOPS SSD (io1) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| IOPS for Provisioned IOPS SSD (io2) volumes | Each supported Region: 100,000 | Yes | The maximum aggregated number of IOPS that can be provisioned across Provisioned IOPS SSD (io2) volumes in this Region. |
| IOPS modifications for Provisioned IOPS SSD (io1) volumes | Each supported Region: 500,000 | Yes | The maximum aggregated number of IOPS that can be requested in volume modifications across Provisioned IOPS SSD (io1) volumes in this Region. |
| IOPS modifications for Provisioned IOPS SSD (io2) volumes | Each supported Region: 100,000 | Yes | The maximum current (from) and requested (to) IOPS for volume modification requests across Provisioned IOPS SSD (io2) volumes in this Region. |
| In-progress snapshot archives per account | Each supported Region: 25 | Yes | The maximum number of in-progress snapshot archives per account. |
| In-progress snapshot restores from archive per account | Each supported Region: 5 | Yes | The maximum number of in-progress snapshot restores from archive per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListChangedBlocks requests per account | Each supported Region: 50 per second | No | The maximum number of ListChangedBlocks requests allowed per account. |
| ListSnapshotBlocks requests per account | Each supported Region: 50 per second | No | The maximum number of ListSnapshotBlocks requests allowed per account. |
| Max concurrent copy volume operations per account | Each supported Region: 5 | Yes | Maximum number of concurrent copy volume operations allowed per account in this region. |
| Provisioned Rate for Volume Initialization across concurrent volume creation requests per Region | Each supported Region: 5,000 | Yes | The maximum aggregated Provisioned Rate for Volume Initialization, in MiB/s, that can be requested across concurrent volume creation requests in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| PutSnapshotBlock requests per account | us-east-1: 5,000 per second<br><br>us-east-2: 5,000 per second<br><br>us-west-2: 5,000 per second<br><br>ap-southeast-1: 5,000 per second<br><br>eu-west-1: 5,000 per second<br><br>Each of the other supported Regions: 1,000 per second | Yes | The maximum number of PutSnapshotBlock requests allowed per account. |
| PutSnapshotBlock requests per snapshot | Each supported Region: 1,000 per second | No | The maximum number of PutSnapshotBlock requests allowed per snapshot. |
| Snapshots per Region | Each supported Region: 100,000 | Yes | The maximum number of snapshots per Region. |
| StartSnapshot pending snapshots per account | Each supported Region: 100 | No | The maximum number of pending snapshots per account that can be created using the StartSnapshot API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StartSnapshot requests per account | Each supported Region: 10 per second | No | The maximum number of StartSnapshot requests allowed per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage for Cold HDD (sc1) volumes, in TiB | af-south-1: 300<br>ap-east-1: 300<br>ap-east-2: 300<br>ap-northeast-3: 300<br>ap-south-2: 300<br>ap-southeast-3: 300<br>ap-southeast-4: 300<br>ap-southeast-5: 300<br>ap-southeast-6: 300<br>ap-southeast-7: 300<br>ca-west-1: 300<br>eu-central-2: 300<br>eu-south-1: 300<br>eu-south-2: 300<br>il-central-1: 300<br>mx-central-1: 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across Cold HDD (sc1) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 50 | | |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage for General Purpose SSD (gp2) volumes, in TiB | af-south-1: 300<br>ap-east-1: 300<br>ap-east-2: 300<br>ap-northeast-3: 300<br>ap-south-2: 300<br>ap-southeast-3: 300<br>ap-southeast-4: 300<br>ap-southeast-5: 300<br>ap-southeast-6: 300<br>ap-southeast-7: 300<br>ca-west-1: 300<br>eu-central-2: 300<br>eu-south-1: 300<br>eu-south-2: 300<br>il-central-1: 300<br>mx-central-1: 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across General Purpose SSD (gp2) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 50 | | |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage for General Purpose SSD (gp3) volumes, in TiB | af-south-1: 300<br>ap-east-1: 300<br>ap-east-2: 300<br>ap-northeast-3: 300<br>ap-south-2: 300<br>ap-southeast-3: 300<br>ap-southeast-4: 300<br>ap-southeast-5: 300<br>ap-southeast-6: 300<br>ap-southeast-7: 300<br>ca-west-1: 300<br>eu-central-2: 300<br>eu-south-1: 300<br>eu-south-2: 300<br>il-central-1: 300<br>mx-central-1: 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across General Purpose SSD (gp3) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 50 | | |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage for Magnetic (standard) volumes, in TiB | af-south-1: 300<br>ap-east-1: 300<br>ap-east-2: 300<br>ap-northeast-3: 300<br>ap-south-2: 300<br>ap-southeast-3: 300<br>ap-southeast-4: 300<br>ap-southeast-5: 300<br>ap-southeast-6: 300<br>ap-southeast-7: 300<br>ca-west-1: 300<br>eu-central-2: 300<br>eu-south-1: 300<br>eu-south-2: 300<br>il-central-1: 300<br>mx-central-1: 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across Magnetic (standard) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 50 | | |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage for Provisioned IOPS SSD (io1) volumes, in TiB | af-south-1: 300<br>ap-east-1: 300<br>ap-east-2: 300<br>ap-northeast-3: 300<br>ap-south-2: 300<br>ap-southeast-3: 300<br>ap-southeast-4: 300<br>ap-southeast-5: 300<br>ap-southeast-6: 300<br>ap-southeast-7: 300<br>ca-west-1: 300<br>eu-central-2: 300<br>eu-south-1: 300<br>eu-south-2: 300<br>il-central-1: 300<br>mx-central-1: 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across Provisioned IOPS SSD (io1) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 50 | | |
| Storage for Provisioned IOPS SSD (io2) volumes, in TiB | Each supported Region: 20 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across Provisioned IOPS SSD (io2) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage for Throughput Optimized HDD (st1) volumes, in TiB | af-south-1: 300<br>ap-east-1: 300<br>ap-east-2: 300<br>ap-northeast-3: 300<br>ap-south-2: 300<br>ap-southeast-3: 300<br>ap-southeast-4: 300<br>ap-southeast-5: 300<br>ap-southeast-6: 300<br>ap-southeast-7: 300<br>ca-west-1: 300<br>eu-central-2: 300<br>eu-south-1: 300<br>eu-south-2: 300<br>il-central-1: 300<br>mx-central-1: 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across Throughput Optimized HDD (st1) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 50 | | |
| Storage modifications for Cold HDD (sc1) volumes, in TiB | Each supported Region: 500 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across Cold HDD (sc1) volumes in this Region. |
| Storage modifications for General Purpose SSD (gp2) volumes, in TiB | Each supported Region: 500 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across General Purpose SSD (gp2) volumes in this Region. |
| Storage modifications for General Purpose SSD (gp3) volumes, in TiB | Each supported Region: 500 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across General Purpose SSD (gp3) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage modifications for Magnetic (standard) volumes, in TiB | Each supported Region: 500 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across Magnetic (standard) volumes in this Region. |
| Storage modifications for Provisioned IOPS SSD (io1) volumes, in TiB | Each supported Region: 500 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across Provisioned IOPS SSD (io1) volumes in this Region. |
| Storage modifications for Provisioned IOPS SSD (io2) volumes, in TiB | Each supported Region: 20 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across Provisioned IOPS SSD (io2) volumes in this Region. |
| Storage modifications for Throughput Optimized HDD (st1) volumes, in TiB | Each supported Region: 500 | Yes | The maximum aggregated amount of storage, in TiB, that can be requested in volume modifications across Throughput Optimized HDD (st1) volumes in this Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Time-based snapshot copy throughput per destination Region | Each supported Region: 2,000 | Yes | The maximum account-level throughput, in MiB/sec, for time-based snapshot copy operations per destination Region. |

# Amazon EC2 endpoints and quotas

**Important**

AWS Regions launched after **October 30, 2021** will no longer support Amazon EC2 API requests over connections that are established using TLSv1, TLSv1.1, or unencrypted HTTP.

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | ec2.us-east-2.amazonaws.com | HTTP and HTTPS | |
| | | ec2-fips.us-east-2.amazonaws.com | | |
| | | ec2.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | | HTTPS | |
| US East (N. Virginia) | us-east-1 | ec2.us-east-1.amazonaws.com<br>ec2-fips.us-east-1.amazonaws.com<br>ec2.us-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | ec2.us-west-1.amazonaws.com<br>ec2-fips.us-west-1.amazonaws.com<br>ec2.us-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | ec2.us-west-2.amazonaws.com<br>ec2-fips.us-west-2.amazonaws.com<br>ec2.us-west-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | ec2.af-south-1.amazonaws.com<br>ec2.af-south-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | ec2.ap-east-1.amazonaws.com<br>ec2.ap-east-1.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | ec2.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ec2.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | ec2.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ec2.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | ec2.ap-south-1.amazonaws.com<br>ec2.ap-south-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ec2.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ec2.ap-northeast-3.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | ec2.ap-northeast-2.amazonaws.com<br>ec2.ap-northeast-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | ec2.ap-southeast-1.amazonaws.com<br>ec2.ap-southeast-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ec2.ap-southeast-2.amazonaws.com<br>ec2.ap-southeast-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | ec2.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | ec2.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | ec2.ap-northeast-1.amazonaws.com<br>ec2.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | ec2.ca-central-1.amazonaws.com<br>ec2-fips.ca-central-1.amazonaws.com<br>ec2.ca-central-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | ec2.ca-west-1.amazonaws.com<br>ec2-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | ec2.eu-central-1.amazonaws.com<br>ec2.eu-central-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | ec2.eu-west-1.amazonaws.com<br>ec2.eu-west-1.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | ec2.eu-west-2.amazonaws.com<br>ec2.eu-west-2.api.aws | HTTP and HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | ec2.eu-south-1.amazonaws.com<br>ec2.eu-south-1.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | ec2.eu-west-3.amazonaws.com<br><br>ec2.eu-west-3.api.aws | HTTP and HTTPS<br><br>HTTPS | |
| Europe (Spain) | eu-south-2 | ec2.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | ec2.eu-north-1.amazonaws.com<br><br>ec2.eu-north-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |
| Europe (Zurich) | eu-central-2 | ec2.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | ec2.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | ec2.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | ec2.me-south-1.amazonaws.com<br><br>ec2.me-south-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |
| Middle East (UAE) | me-central-1 | ec2.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | ec2.sa-east-1.amazonaws.com<br><br>ec2.sa-east-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | ec2.us-gov-east-1.amazonaws.com<br><br>ec2.us-gov-east-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | ec2.us-gov-west-1.amazonaws.com<br><br>ec2.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

If you specify the general endpoint (ec2.amazonaws.com), Amazon EC2 directs your request to the endpoint for `us-east-1`.

## Service quotas

### Amazon EC2

The following are the quotas for Amazon EC2.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amazon FPGA images (AFIs) | 100 | Yes | The maximum number of available Amazon FPGA images (AFIs) that you can own in this Region. |
| AMI sharing | 1,000 | Yes | The maximum number of entities (organizations, organizational units, and accounts) that an AMI can be shared with in this Region. Note that if you |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | share an AMI with an organization, the number of accounts in the organization does not count towards the quota. |
| AMIs | 50,000 | Yes | The maximum number of public and private AMIs allowed in this Region. These include available, disabled, and pending AMIs, and AMIs in the Recycle Bin. |
| EC2-VPC Elastic IPs | 5 | Yes | The maximum number of Elastic IP addresses that you can allocate for EC2-VPC in this Region. |
| Launch template versions | 10,000 | No | Maximum number of launch template versions per launch template per Region per account. |
| Launch templates | 5,000 | No | Maximum number of launch templates per Region per account.<br><br>For more information, see Launch template restrictions. |
| New Reserved Instances per month | 20 | Yes | The maximum number of Reserved Instances (RIs) that you can purchase per month in the current account. For regional RIs, this is the maximum number of RIs that you can purchase for the current Region. For zonal RIs, this is the maximum number of RIs that you can purchase for each Availability Zone in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Public AMIs | 5 | Yes | The maximum number of public AMIs, including public AMIs in the Recycle Bin, allowed in this Region. |
| Secondary Networks per Region | 5 | Yes | The maximum number of secondary networks per Region |
| Subnets per Secondary Network | 200 | Yes | The maximum number of subnets per secondary network |

### Instance types

For details, see Quotas in the *Amazon EC2 Instance Types* guide.

### EC2 Fast Launch

The following quotas are for EC2 Fast Launch.

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Parallel instance launches | Each supported Region: 40 | Yes | A setting that you configure for the EC2 Fast Launch feature that defines the maximum number of Windows instances it can launch at the same time to create pre-provisioned snapshots in the owner's account in the current Region. For more information, see https://docs.aws.amazon.com/AWSEC2/latest/Win |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | dowsGuide/win-ami-config-fast-launch.html |

### VM Import/Export

The following quotas are for VM Import/Export.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrent task limit for ImportImage, ImportSnapshot, and ExportImage | Each supported Region: 20 | Yes | The maximum number of concurrent tasks for a given account initiated by the following VM Import/Export APIs: ImportImage, ImportSnapshot, and ExportImage. |
| Concurrent task limit for ImportInstance, ImportVolume, and CreateInstanceExportTask | Each supported Region: 5 | Yes | The maximum number of concurrent tasks for a given account initiated by the following VM Import/Export APIs: ImportInstance, ImportVolume, and CreateInstanceExportTask. |

## Amazon EC2 Auto Scaling endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | autoscaling.us-east-2.amazonaws.com<br>autoscaling-fips.us-east-2.api.aws<br>autoscaling-fips.us-east-2.amazonaws.com<br>autoscaling.us-east-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US East (N. Virginia) | us-east-1 | autoscaling.us-east-1.amazonaws.com<br>autoscaling-fips.us-east-1.api.aws<br>autoscaling-fips.us-east-1.amazonaws.com<br>autoscaling.us-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (N. | us-west-1 | autoscaling.us-west-1.amazonaws.com<br>autoscaling-fips.us-west-1.api.aws | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| California) | | autoscaling-fips.us-west-1.amazonaws.com<br>autoscaling.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | autoscaling.us-west-2.amazonaws.com<br>autoscaling-fips.us-west-2.api.aws<br>autoscaling-fips.us-west-2.amazonaws.com<br>autoscaling.us-west-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| Africa (Cape Town) | af-south-1 | autoscaling.af-south-1.amazonaws.com<br>autoscaling.af-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | autoscaling.ap-east-1.amazonaws.com<br>autoscaling.ap-east-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | autoscaling.ap-south-2.amazonaws.com<br>autoscaling.ap-south-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | autoscaling.ap-southeast-3.amazonaws.com<br>autoscaling.ap-southeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | autoscaling.ap-southeast-5.amazonaws.com<br>autoscaling.ap-southeast-5.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | autoscaling.ap-southeast-4.amazonaws.com<br>autoscaling.ap-southeast-4.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | autoscaling.ap-south-1.amazonaws.com<br>autoscaling.ap-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | autoscaling.ap-southeast-6.amazonaws.com<br>autoscaling.ap-southeast-6.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | autoscaling.ap-northeast-3.amazonaws.com<br>autoscaling.ap-northeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | autoscaling.ap-northeast-2.amazonaws.com<br>autoscaling.ap-northeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | autoscaling.ap-southeast-1.amazonaws.com<br>autoscaling.ap-southeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | autoscaling.ap-southeast-2.amazonaws.com<br>autoscaling.ap-southeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | autoscaling.ap-east-2.amazonaws.com<br>autoscaling.ap-east-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | autoscaling.ap-southeast-7.amazonaws.com<br>autoscaling.ap-southeast-7.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | autoscaling.ap-northeast-1.amazonaws.com<br>autoscaling.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | autoscaling.ca-central-1.amazonaws.com<br>autoscaling-fips.ca-central-1.api.aws<br>autoscaling-fips.ca-central-1.amazonaws.com<br>autoscaling.ca-central-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| Canada West (Calgary) | ca-west-1 | autoscaling.ca-west-1.amazonaws.com<br>autoscaling-fips.ca-west-1.api.aws<br>autoscaling-fips.ca-west-1.amazonaws.com<br>autoscaling.ca-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | autoscaling.eu-central-1.amazonaws.com<br>autoscaling.eu-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | autoscaling.eu-west-1.amazonaws.com<br>autoscaling.eu-west-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (London) | eu-west-2 | autoscaling.eu-west-2.amazonaws.com<br>autoscaling.eu-west-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Milan) | eu-south-1 | autoscaling.eu-south-1.amazonaws.com<br>autoscaling.eu-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | autoscaling.eu-west-3.amazonaws.com<br><br>autoscaling.eu-west-3.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | autoscaling.eu-south-2.amazonaws.com<br><br>autoscaling.eu-south-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | autoscaling.eu-north-1.amazonaws.com<br><br>autoscaling.eu-north-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Zurich) | eu-central-2 | autoscaling.eu-central-2.amazonaws.com<br><br>autoscaling.eu-central-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | autoscaling.il-central-1.amazonaws.com<br>autoscaling.il-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Mexico (Central) | mx-central-1 | autoscaling.mx-central-1.amazonaws.com<br>autoscaling.mx-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | autoscaling.me-south-1.amazonaws.com<br>autoscaling.me-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Middle East (UAE) | me-central-1 | autoscaling.me-central-1.amazonaws.com<br>autoscaling.me-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | autoscaling.sa-east-1.amazonaws.com<br><br>autoscaling.sa-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | autoscaling.us-gov-east-1.amazonaws.com<br><br>autoscaling.us-gov-east-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | autoscaling.us-gov-west-1.amazonaws.com<br><br>autoscaling.us-gov-west-1.api.aws | HTTP and HTTPS<br><br>HTTPS | |

If you specify the general endpoint (autoscaling.amazonaws.com), Amazon EC2 Auto Scaling directs your request to the endpoint for `us-east-1`.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Auto Scaling groups per region | Each supported Region: 500 | Yes | The maximum number of Auto Scaling groups allowed for your AWS account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Classic Load Balancers per Auto Scaling group | Each supported Region: 50 | No | The maximum number of Classic Load Balancers per Auto Scaling group |
| Launch configurations per region | Each supported Region: 200 | Yes | The maximum number of launch configurations allowed for your AWS account |
| Lifecycle hooks per Auto Scaling group | Each supported Region: 50 | No | The maximum number of lifecycle hooks per Auto Scaling group |
| SNS topics per Auto Scaling group | Each supported Region: 10 | No | The maximum number of SNS topics per Auto Scaling group |
| Scaling policies per Auto Scaling group | Each supported Region: 50 | No | The maximum number of scaling policies per Auto Scaling group |
| Scheduled actions per Auto Scaling group | Each supported Region: 125 | No | The maximum number of scheduled actions per Auto Scaling group |
| Step adjustments per step scaling policy | Each supported Region: 20 | No | The maximum number of step adjustments per step scaling policy |
| Target groups per Auto Scaling group | Each supported Region: 50 | No | The maximum number of target groups per Auto Scaling group |

For more information, see Quotas for Amazon EC2 Auto Scaling in the *Amazon EC2 Auto Scaling User Guide*.

# EC2 Image Builder endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | imagebuilder.us-east-2.amazonaws.com | HTTPS | |
| | | imagebuilder.us-east-2.api.aws | HTTPS | |
| | | imagebuilder-fips.us-east-2.amazonaws.com | HTTPS | |
| | | imagebuilder-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | imagebuilder.us-east-1.amazonaws.com | HTTPS | |
| | | imagebuilder.us-east-1.api.aws | HTTPS | |
| | | imagebuilder-fips.us-east-1.amazonaws.com | HTTPS | |
| | | imagebuilder-fips.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | imagebuilder.us-west-1.amazonaws.com | HTTPS | |
| | | imagebuilder.us-west-1.api.aws | HTTPS | |
| | | imagebuilder-fips.us-west-1.amazonaws.com | HTTPS | |
| | | imagebuilder-fips.us-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | imagebuilder.us-west-2.amazonaws.com | HTTPS |
| | | imagebuilder.us-west-2.api.aws | HTTPS |
| | | imagebuilder-fips.us-west-2.amazonaws.com | HTTPS |
| | | imagebuilder-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | imagebuilder.af-south-1.amazonaws.com | HTTPS |
| | | imagebuilder.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | imagebuilder.ap-east-1.amazonaws.com | HTTPS |
| | | imagebuilder.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | imagebuilder.ap-south-2.amazonaws.com | HTTPS |
| | | imagebuilder.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | imagebuilder.ap-southeast-3.amazonaws.com | HTTPS |
| | | imagebuilder.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | imagebuilder.ap-southeast-5.amazonaws.com | HTTPS |
| | | imagebuilder.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | imagebuilder.ap-southeast-4.amazonaws.com | HTTPS |
| | | imagebuilder.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | imagebuilder.ap-south-1.amazonaws.com<br>imagebuilder.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | imagebuilder.ap-southeast-6.amazonaws.com<br>imagebuilder.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | imagebuilder.ap-northeast-3.amazonaws.com<br>imagebuilder.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | imagebuilder.ap-northeast-2.amazonaws.com<br>imagebuilder.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | imagebuilder.ap-southeast-1.amazonaws.com<br>imagebuilder.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | imagebuilder.ap-southeast-2.amazonaws.com<br>imagebuilder.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | imagebuilder.ap-east-2.amazonaws.com<br>imagebuilder.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | imagebuilder.ap-southeast-7.amazonaws.com<br>imagebuilder.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | imagebuilder.ap-northeast-1.amazonaws.com | HTTPS |
| | | imagebuilder.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | imagebuilder.ca-central-1.amazonaws.com | HTTPS |
| | | imagebuilder.ca-central-1.api.aws | HTTPS |
| | | imagebuilder-fips.ca-central-1.amazonaws.com | HTTPS |
| | | imagebuilder-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | imagebuilder.ca-west-1.amazonaws.com | HTTPS |
| | | imagebuilder.ca-west-1.api.aws | HTTPS |
| | | imagebuilder-fips.ca-west-1.amazonaws.com | HTTPS |
| | | imagebuilder-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | imagebuilder.eu-central-1.amazonaws.com | HTTPS |
| | | imagebuilder.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | imagebuilder.eu-west-1.amazonaws.com | HTTPS |
| | | imagebuilder.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | imagebuilder.eu-west-2.amazonaws.com | HTTPS |
| | | imagebuilder.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | imagebuilder.eu-south-1.amazonaws.com | HTTPS |
| | | imagebuilder.eu-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | imagebuilder.eu-west-3.amazonaws.com<br>imagebuilder.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | imagebuilder.eu-south-2.amazonaws.com<br>imagebuilder.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | imagebuilder.eu-north-1.amazonaws.com<br>imagebuilder.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | imagebuilder.eu-central-2.amazonaws.com<br>imagebuilder.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | imagebuilder.il-central-1.amazonaws.com<br>imagebuilder.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | imagebuilder.mx-central-1.amazonaws.com<br>imagebuilder.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | imagebuilder.me-south-1.amazonaws.com<br>imagebuilder.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | imagebuilder.me-central-1.amazonaws.com<br>imagebuilder.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | imagebuilder.sa-east-1.amazonaws.com<br>imagebuilder.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | imagebuilder.us-gov-east-1.amazonaws.com<br>imagebuilder.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | imagebuilder.us-gov-west-1.amazonaws.com<br>imagebuilder.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Component parameter length | Each supported Region: 1,024 | Yes | The maximum number of characters that a component parameter value can contain. |
| Component size | Each supported Region: 64 Kilobytes | Yes | The maximum size of the data field of a component. |
| Components | Each supported Region: 1,000 | Yes | The maximum number of components that you can create in an account in the current Region. |
| Components per image recipe | Each supported Region: 20 | Yes | The maximum number of components that can be associated with a single image recipe. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrent AMI copies per distribution configuration | Each supported Region: 50 | Yes | The maximum number of target accounts that can be defined for a single distribution configuration. |
| Concurrent builds | Each supported Region: 100 | Yes | The maximum number of concurrent builds that can be in progress in this account in the current Region. |
| Concurrent resource state update | Each supported Region: 50 | Yes | The maximum number of concurrent resource state updates that can be in progress in this account in the current Region. |
| Container recipes | Each supported Region: 1,000 | Yes | The maximum number of container recipes that you can create in an account in the current Region. |
| Distribution configurations | Each supported Region: 1,000 | Yes | The maximum number of distribution configurations that you can create in an account in the current Region. |
| Docker template size | Each supported Region: 64 Kilobytes | Yes | The maximum size of a Docker template data field. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Image pipelines | Each supported Region: 75 | Yes | The maximum number of image pipelines that you can create in an account in the current Region. |
| Image recipes | Each supported Region: 1,000 | Yes | The maximum number of image recipes that you can create in an account in the current Region. |
| Image workflow parameter | Each supported Region: 1,024 | Yes | The maximum number of characters that an image workflow parameter value can contain. |
| Image workflow size | Each supported Region: 64 Kilobytes | Yes | The maximum size of the data field of an image workflow. |
| Image workflows | Each supported Region: 1,000 | Yes | The maximum number of image workflows that you can create in an account in the current Region. |
| Image workflows per image | Each supported Region: 10 | Yes | The maximum number of image workflows that can be associated with a single image. |
| Infrastructure configurations | Each supported Region: 1,000 | Yes | The maximum number of infrastructure configurations that you can create in an account in the current Region. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Launch templates modified per distribution configuration Region | Each supported Region: 5 | Yes | The maximum number of launch templates that a single distribution configuration can modify per Region. |
| Lifecycle policies | Each supported Region: 50 | Yes | The maximum number of lifecycle policies that you can create in an account in the current Region. |
| Parameters per component | Each supported Region: 25 | Yes | The maximum number of parameters that a single component can contain. |
| Parameters per image workflow | Each supported Region: 25 | Yes | The maximum number of parameters that a single image workflow can contain. |
| SSM Parameters updated per distribution configuration Region | Each supported Region: 10 | Yes | The maximum number of SSM Parameters that a single distribution configuration can update per Region. |
| Steps per image workflow | Each supported Region: 15 | Yes | The maximum number of steps that a single image workflow can contain. |

# Amazon EC2 Instance Connect endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | ec2-instance-connect.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | ec2-instance-connect.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | ec2-instance-connect.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | ec2-instance-connect.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | ec2-instance-connect.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | ec2-instance-connect.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific | ap-south-2 | ec2-instance-connect.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ec2-instance-connect.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | ec2-instance-connect.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ec2-instance-connect.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | ec2-instance-connect.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ec2-instance-connect.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ec2-instance-connect.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ec2-instance-connect.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | ec2-instance-connect.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ec2-instance-connect.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | ec2-instance-connect.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | ec2-instance-connect.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | ec2-instance-connect.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | ec2-instance-connect.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | ec2-instance-connect.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | ec2-instance-connect.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | ec2-instance-connect.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | ec2-instance-connect.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | ec2-instance-connect.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | ec2-instance-connect.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | ec2-instance-connect.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | ec2-instance-connect.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | ec2-instance-connect.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | ec2-instance-connect.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | ec2-instance-connect.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | ec2-instance-connect.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | ec2-instance-connect.me-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | ec2-instance-connect.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | ec2-instance-connect.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | ec2-instance-connect.us-gov-west-1.amazonaws.com | HTTPS | |

### Service quotas

EC2 Instance Connect has no service quotas.

> **Note**
>
> EC2 Instance Connect should not be confused with EC2 Instance Connect Endpoint, which is a separate service with its own quotas. For more information, see Quotas for EC2 Instance Connect Endpoint in the *Amazon EC2 User Guide*.

## Amazon ECR endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The `ecr` and `api.ecr` endpoints are used for calls to the Amazon ECR API. API actions such as `DescribeImages` and `CreateRepository` go to this endpoint. While the two endpoints function the same, the `api.ecr` endpoint is recommended and the default when using the AWS CLI or AWS SDKs. When connecting to Amazon ECR through an AWS PrivateLink VPC endpoint, you must use the `api.ecr` endpoint to make API calls. For more information, see Amazon ECR Interface VPC Endpoints (AWS PrivateLink) in the *Amazon Elastic Container Registry User Guide*.

For more information about FIPS endpoints, see FIPS endpoints.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | ecr.us-east-2.amazonaws.com | HTTPS | |
| | | ecr-fips.us-east-2.api.aws | HTTPS | |
| | | dkr-ecr.us-east-2.on.aws | HTTPS | |
| | | dkr.ecr.us-east-2.amazonaws.com | HTTPS | |
| | | ecr.us-east-2.api.aws | HTTPS | |
| | | dkr-ecr-fips.us-east-2.on.aws | HTTPS | |
| | | ecr-fips.us-east-2.amazonaws.com | HTTPS | |
| | | api.ecr-fips.us-east-2.amazonaws.com | HTTPS | |
| | | api.ecr.us-east-2.amazonaws.com | HTTPS | |
| | | dkr.ecr-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | ecr.us-east-1.amazonaws.com | HTTPS | |
| | | dkr.ecr-fips.us-east-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | dkr.ecr.us-east-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.us-east-1.on.aws | HTTPS | |
| | | ecr.us-east-1.api.aws | HTTPS | |
| | | dkr-ecr-fips.us-east-1.on.aws | HTTPS | |
| | | api.ecr.us-east-1.amazonaws.com | HTTPS | |
| | | api.ecr-fips.us-east-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | ecr.us-west-1.amazonaws.com | HTTPS | |
| | | api.ecr-fips.us-west-1.amazonaws.com | HTTPS | |
| | | ecr.us-west-1.api.aws | HTTPS | |
| | | ecr-fips.us-west-1.api.aws | HTTPS | |
| | | dkr-ecr.us-west-1.on.aws | HTTPS | |
| | | dkr-ecr-fips.us-west-1.on.aws | HTTPS | |
| | | dkr.ecr.us-west-1.amazonaws.com | HTTPS | |
| | | dkr.ecr-fips.us-west-1.amazonaws.com | HTTPS | |
| | | api.ecr.us-west-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | ecr.us-west-2.amazonaws.com | HTTPS | |
| | | ecr-fips.us-west-2.api.aws | HTTPS | |
| | | ecr.us-west-2.api.aws | HTTPS | |
| | | ecr-fips.us-west-2.amazonaws.com | HTTPS | |
| | | dkr.ecr.us-west-2.amazonaws.com | HTTPS | |
| | | dkr.ecr-fips.us-west-2.amazonaws.com | HTTPS | |
| | | api.ecr-fips.us-west-2.amazonaws.com | HTTPS | |
| | | dkr-ecr.us-west-2.on.aws | HTTPS | |
| | | dkr-ecr-fips.us-west-2.on.aws | HTTPS | |
| | | api.ecr.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | ecr.af-south-1.amazonaws.com | HTTPS | |
| | | api.ecr.af-south-1.amazonaws.com | HTTPS | |
| | | ecr.af-south-1.api.aws | HTTPS | |
| | | dkr.ecr.af-south-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.af-south-1.on.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | ecr.ap-east-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-east-1.on.aws | HTTPS | |
| | | dkr.ecr.ap-east-1.amazonaws.com | HTTPS | |
| | | ecr.ap-east-1.api.aws | HTTPS | |
| | | api.ecr.ap-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | ecr.ap-south-2.amazonaws.com | HTTPS | |
| | | api.ecr.ap-south-2.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-south-2.on.aws | HTTPS | |
| | | dkr.ecr.ap-south-2.amazonaws.com | HTTPS | |
| | | ecr.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ecr.ap-southeast-3.amazonaws.com | HTTPS | |
| | | ecr.ap-southeast-3.api.aws | HTTPS | |
| | | api.ecr.ap-southeast-3.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-southeast-3.on.aws | HTTPS | |
| | | dkr.ecr.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | ecr.ap-southeast-5.amazonaws.com | HTTPS | |
| | | api.ecr.ap-southeast-5.amazonaws.com | HTTPS | |
| | | dkr.ecr.ap-southeast-5.amazonaws.com | HTTPS | |
| | | ecr.ap-southeast-5.api.aws | HTTPS | |
| | | dkr-ecr.ap-southeast-5.on.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ecr.ap-southeast-4.amazonaws.com | HTTPS | |
| | | api.ecr.ap-southeast-4.amazonaws.com | HTTPS | |
| | | ecr.ap-southeast-4.api.aws | HTTPS | |
| | | dkr-ecr.ap-southeast-4.on.aws | HTTPS | |
| | | dkr.ecr.ap-southeast-4.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | ecr.ap-south-1.amazonaws.com | HTTPS | |
| | | api.ecr.ap-south-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-south-1.on.aws | HTTPS | |
| | | dkr.ecr.ap-south-1.amazonaws.com | HTTPS | |
| | | ecr.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ecr.ap-southeast-6.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-southeast-6.on.aws | HTTPS | |
| | | ecr.ap-southeast-6.api.aws | HTTPS | |
| | | api.ecr.ap-southeast-6.amazonaws.com | HTTPS | |
| | | dkr.ecr.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ecr.ap-northeast-3.amazonaws.com | HTTPS | |
| | | api.ecr.ap-northeast-3.amazonaws.com | HTTPS | |
| | | dkr.ecr.ap-northeast-3.amazonaws.com | HTTPS | |
| | | ecr.ap-northeast-3.api.aws | HTTPS | |
| | | dkr-ecr.ap-northeast-3.on.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ecr.ap-northeast-2.amazonaws.com | HTTPS | |
| | | api.ecr.ap-northeast-2.amazonaws.com | HTTPS | |
| | | ecr.ap-northeast-2.api.aws | HTTPS | |
| | | dkr.ecr.ap-northeast-2.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-northeast-2.on.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | ecr.ap-southeast-1.amazonaws.com | HTTPS |
| | | api.ecr.ap-southeast-1.amazonaws.com | HTTPS |
| | | dkr.ecr.ap-southeast-1.amazonaws.com | HTTPS |
| | | ecr.ap-southeast-1.api.aws | HTTPS |
| | | dkr-ecr.ap-southeast-1.on.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ecr.ap-southeast-2.amazonaws.com | HTTPS |
| | | ecr.ap-southeast-2.api.aws | HTTPS |
| | | api.ecr.ap-southeast-2.amazonaws.com | HTTPS |
| | | dkr.ecr.ap-southeast-2.amazonaws.com | HTTPS |
| | | dkr-ecr.ap-southeast-2.on.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | ecr.ap-east-2.amazonaws.com | HTTPS |
| | | api.ecr.ap-east-2.amazonaws.com | HTTPS |
| | | dkr.ecr.ap-east-2.amazonaws.com | HTTPS |
| | | dkr-ecr.ap-east-2.on.aws | HTTPS |
| | | ecr.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | ecr.ap-southeast-7.amazonaws.com | HTTPS |
| | | dkr-ecr.ap-southeast-7.on.aws | HTTPS |
| | | dkr.ecr.ap-southeast-7.amazonaws.com | HTTPS |
| | | ecr.ap-southeast-7.api.aws | HTTPS |
| | | api.ecr.ap-southeast-7.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | ecr.ap-northeast-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.ap-northeast-1.on.aws | HTTPS | |
| | | api.ecr.ap-northeast-1.amazonaws.com | HTTPS | |
| | | ecr.ap-northeast-1.api.aws | HTTPS | |
| | | dkr.ecr.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | ecr.ca-central-1.amazonaws.com | HTTPS | |
| | | ecr.ca-central-1.api.aws | HTTPS | |
| | | dkr-ecr.ca-central-1.on.aws | HTTPS | |
| | | dkr.ecr.ca-central-1.amazonaws.com | HTTPS | |
| | | api.ecr.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | ecr.ca-west-1.amazonaws.com | HTTPS | |
| | | ecr.ca-west-1.api.aws | HTTPS | |
| | | api.ecr.ca-west-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.ca-west-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.ca-west-1.on.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | ecr.eu-central-1.amazonaws.com | HTTPS | |
| | | ecr.eu-central-1.api.aws | HTTPS | |
| | | dkr-ecr.eu-central-1.on.aws | HTTPS | |
| | | api.ecr.eu-central-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | ecr.eu-west-1.amazonaws.com | HTTPS | |
| | | api.ecr.eu-west-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.eu-west-1.on.aws | HTTPS | |
| | | dkr.ecr.eu-west-1.amazonaws.com | HTTPS | |
| | | ecr.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | ecr.eu-west-2.amazonaws.com | HTTPS | |
| | | ecr.eu-west-2.api.aws | HTTPS | |
| | | api.ecr.eu-west-2.amazonaws.com | HTTPS | |
| | | dkr.ecr.eu-west-2.amazonaws.com | HTTPS | |
| | | dkr-ecr.eu-west-2.on.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | ecr.eu-south-1.amazonaws.com | HTTPS | |
| | | ecr.eu-south-1.api.aws | HTTPS | |
| | | api.ecr.eu-south-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.eu-south-1.on.aws | HTTPS | |
| | | dkr.ecr.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | ecr.eu-west-3.amazonaws.com | HTTPS | |
| | | ecr.eu-west-3.api.aws | HTTPS | |
| | | api.ecr.eu-west-3.amazonaws.com | HTTPS | |
| | | dkr.ecr.eu-west-3.amazonaws.com | HTTPS | |
| | | dkr-ecr.eu-west-3.on.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | ecr.eu-south-2.amazonaws.com | HTTPS | |
| | | ecr.eu-south-2.api.aws | HTTPS | |
| | | dkr-ecr.eu-south-2.on.aws | HTTPS | |
| | | api.ecr.eu-south-2.amazonaws.com | HTTPS | |
| | | dkr.ecr.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | ecr.eu-north-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.eu-north-1.amazonaws.com | HTTPS | |
| | | api.ecr.eu-north-1.amazonaws.com | HTTPS | |
| | | ecr.eu-north-1.api.aws | HTTPS | |
| | | dkr-ecr.eu-north-1.on.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | ecr.eu-central-2.amazonaws.com | HTTPS | |
| | | ecr.eu-central-2.api.aws | HTTPS | |
| | | dkr-ecr.eu-central-2.on.aws | HTTPS | |
| | | api.ecr.eu-central-2.amazonaws.com | HTTPS | |
| | | dkr.ecr.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | ecr.il-central-1.amazonaws.com | HTTPS | |
| | | ecr.il-central-1.api.aws | HTTPS | |
| | | dkr-ecr.il-central-1.on.aws | HTTPS | |
| | | api.ecr.il-central-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | ecr.mx-central-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.mx-central-1.amazonaws.com | HTTPS | |
| | | api.ecr.mx-central-1.amazonaws.com | HTTPS | |
| | | ecr.mx-central-1.api.aws | HTTPS | |
| | | dkr-ecr.mx-central-1.on.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | ecr.me-south-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.me-south-1.amazonaws.com | HTTPS | |
| | | ecr.me-south-1.api.aws | HTTPS | |
| | | api.ecr.me-south-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.me-south-1.on.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | ecr.me-central-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.me-central-1.amazonaws.com | HTTPS | |
| | | ecr.me-central-1.api.aws | HTTPS | |
| | | api.ecr.me-central-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.me-central-1.on.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | ecr.sa-east-1.amazonaws.com | HTTPS | |
| | | dkr-ecr.sa-east-1.on.aws | HTTPS | |
| | | dkr.ecr.sa-east-1.amazonaws.com | HTTPS | |
| | | ecr.sa-east-1.api.aws | HTTPS | |
| | | api.ecr.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | ecr.us-gov-east-1.amazonaws.com | HTTPS | |
| | | ecr.us-gov-east-1.api.aws | HTTPS | |
| | | dkr.ecr.us-gov-east-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | api.ecr-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | api.ecr.us-gov-east-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-gov-east-1.api.aws | HTTPS | |
| | | dkr.ecr-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | dkr-ecr-fips.us-gov-east-1.on.aws | HTTPS | |
| | | dkr-ecr.us-gov-east-1.on.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | ecr.us-gov-west-1.amazonaws.com | HTTPS | |
| | | dkr.ecr-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | api.ecr.us-gov-west-1.amazonaws.com | HTTPS | |
| | | dkr.ecr.us-gov-west-1.amazonaws.com | HTTPS | |
| | | api.ecr-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | ecr-fips.us-gov-west-1.api.aws | HTTPS | |
| | | ecr.us-gov-west-1.api.aws | HTTPS | |
| | | dkr-ecr-fips.us-gov-west-1.on.aws | HTTPS | |
| | | dkr-ecr.us-gov-west-1.on.aws | HTTPS | |

Additionally, the Amazon ECR API has dual-stack endpoints. The dual-stack endpoint was introduced in April 2025. To use the dual-stack endpoints with the AWS CLI, see the Dual-stack and FIPS endpoints configuration in the *AWS SDKs and Tools Reference Guide*. The following lists the new endpoints:

- OCI and Docker client requests:

  `<registry-id>.dkr-ecr.<aws-region>.on.aws`

- AWS API requests:

  `ecr.<aws-region>.api.aws`

## Service quotas

The following table provides the default limits for Amazon Elastic Container Registry (Amazon ECR).

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Basic image scans per 24 hours | Each supported Region: 100,000 | No | The maximum number of images that can be scanned within a 24 hour period in the current account and region using basic scanning. This limit includes both scan on push and manual scans. |
| Filters per rule in a replication configuration | Each supported Region: 100 | No | The maximum number of filters per rule in a replication configuration. |
| Images per repository | Each supported Region: 100,000 | Yes | The maximum number of images per repository. |
| Layer parts | Each supported Region: 4,200 | No | The maximum number of layer parts. This is |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations. |
| Lifecycle policy length | Each supported Region: 30,720 | No | The maximum number of characters in a lifecycle policy. |
| Maximum layer part size | Each supported Region: 10 | No | The maximum size (MiB) of a layer part. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations. |
| Maximum layer size | Each supported Region: 52,000 | No | The maximum size (MiB) of a layer. |
| Minimum layer part size | Each supported Region: 5 | No | The minimum size (MiB) of a layer part. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations. |
| Pull through cache rules per registry | Each supported Region: 50 | No | The maximum number of pull-through cache rules. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of BatchCheckLayerAvailability requests | Each supported Region: 1,000 per second | Yes | The maximum number of BatchCheckLayerAvailability requests that you can make per second in the current Region. When an image is pushed to a repository, each image layer is checked to verify if it has been uploaded before. If it has been uploaded, then the image layer is skipped. |
| Rate of BatchGetImage requests | Each supported Region: 2,000 per second | Yes | The maximum number of BatchGetImage requests that you can make per second in the current Region. When an image is pulled, the BatchGetImage API is called once to retrieve the image manifest. If you request a quota increase for this API, review your GetDownloadUrlForLayer usage as well. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CompleteLayerUpload requests | Each supported Region: 100 per second | Yes | The maximum number of CompleteLayerUpload requests that you can make per second in the current Region. When an image is pushed, the CompleteLayerUpload API is called once per each new image layer to verify that the upload has completed. |
| Rate of GetAuthorizationToken requests | Each supported Region: 500 per second | Yes | The maximum number of GetAuthorizationToken requests that you can make per second in the current Region. |
| Rate of GetDownloadUrlForLayer requests | Each supported Region: 3,000 per second | Yes | The maximum number of GetDownloadUrlForLayer requests that you can make per second in the current Region. When an image is pulled, the GetDownloadUrlForLayer API is called once per image layer that is not already cached. If you request a quota increase for this API, review your BatchGetImage usage as well. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of InitiateLayerUpload requests | Each supported Region: 100 per second | Yes | The maximum number of InitiateLayerUpload requests that you can make per second in the current Region. When an image is pushed, the InitiateLayerUpload API is called once per image layer that has not already been uploaded. Whether or not an image layer has been uploaded is determined by the BatchCheckLayerAvailability API action. |
| Rate of PutImage requests | Each supported Region: 10 per second | Yes | The maximum number of PutImage requests that you can make per second in the current Region. When an image is pushed and all new image layers have been uploaded, the PutImage API is called once to create or update the image manifest and the tags associated with the image. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UploadLayerPart requests | Each supported Region: 500 per second | Yes | The maximum number of UploadLayerPart requests that you can make per second in the current Region. When an image is pushed, each new image layer is uploaded in parts and the UploadLayerPart API is called once per each new image layer part. |
| Rate of image scans | Each supported Region: 1 | No | The maximum number of image scans per image, per 24 hours. |
| Registered repositories | Each supported Region: 100,000 | Yes | The maximum number of repositories that you can create in this account in the current Region. |
| Rules per lifecycle policy | Each supported Region: 50 | No | The maximum number of rules in a lifecycle policy |
| Rules per replication configuration | Each supported Region: 25 | No | The maximum number of rules in a replication configuration. |
| Tags per image | Each supported Region: 1,000 | No | The maximum number of tags per image. |
| Unique destinations across all rules in a replication configuration | Each supported Region: 25 | No | The maximum number of unique destinations across all rules in a replication configuration. |

For more information, see Amazon ECR Service Quotas in the *Amazon Elastic Container Registry User Guide*.

# Amazon ECR Public endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The `ecr-public` and `api.ecr-public` endpoints are used for calls to the Amazon ECR Public API. API actions such as `DescribeImages` and `CreateRepository` go to this endpoint. While the two endpoints function the same, the `api.ecr-public` endpoint is recommended and the default when using the AWS CLI or AWS SDKs.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ecr-public.us-east-1.amazonaws.com | HTTPS | |
| | | ecr-public.us-east-1.api.aws | HTTPS | |
| | | api.ecr-public.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | ecr-public.us-west-2.amazonaws.com | HTTPS | |
| | | api.ecr-public.us-west-2.amazonaws.com | HTTPS | |

## Service quotas

The following are the service quotas for Amazon ECR Public.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Images per repository | Each supported Region: 100,000 | Yes | The maximum number of images per repository. |
| Layer parts | Each supported Region: 1,000 | No | The maximum number of layer parts. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations. |
| Maximum layer part size | Each supported Region: 10 | No | The maximum size (MiB) of a layer part. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations. |
| Maximum layer size | Each supported Region: 10,000 | No | The maximum size per layer. |
| Minimum layer part size | Each supported Region: 5 | No | The minimum size (MiB) of a layer part. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations. |
| Rate of BatchCheckLayerAvailability requests | Each supported Region: 200 per second | Yes | The maximum number of BatchCheckLayerAvailability requests that you can make per second in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | When an image is pushed to a repository, each image layer is checked to verify if it has been uploaded before. If it has been uploaded, then the image layer is skipped. |
| Rate of CompleteLayerUpload requests | Each supported Region: 10 per second | Yes | The maximum number of CompleteLayerUpload requests that you can make per second in the current Region. When an image is pushed, the CompleteLayerUpload API is called once per each new image layer to verify that the upload has completed. |
| Rate of GetAuthorizationToken requests | Each supported Region: 200 per second | Yes | The maximum number of GetAuthorizationToken requests that you can make per second in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of InitiateLayerUpload requests | Each supported Region: 10 per second | Yes | The maximum number of InitiateLayerUpload requests that you can make per second in the current Region. When an image is pushed, the InitiateLayerUpload API is called once per image layer that has not already been uploaded. Whether or not an image layer has been uploaded is determined by the BatchCheckLayerAvailability API action. |
| Rate of PutImage requests | Each supported Region: 10 per second | Yes | The maximum number of PutImage requests that you can make per second in the current Region. When an image is pushed and all new image layers have been uploaded, the PutImage API is called once to create or update the image manifest and the tags associated with the image. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UploadLayerPart requests | Each supported Region: 260 per second | Yes | The maximum number of UploadLayerPart requests that you can make per second in the current Region. When an image is pushed, each new image layer is uploaded in parts and the UploadLayerPart API is called once per each new image layer part. |
| Rate of authenticated image pulls | Each supported Region: 10 per second | Yes | The maximum number of authenticated image pulls per second. |
| Rate of image pulls to AWS resources | Each supported Region: 10 per second | No | The maximum number of image pulls per second to resources running on Amazon ECS, Fargate, or Amazon EC2. |
| Rate of unauthenticated image pulls | Each supported Region: 1 per second | No | The maximum number of unauthenticated image pulls per second. |
| Registered repositories | Each supported Region: 10,000 | Yes | The maximum number of repositories that you can create in this account in the current Region. |
| Tags per image | Each supported Region: 1,000 | No | The maximum number of tags per image. |

For more information, see Amazon ECR Public service quotas in the *Amazon ECR Public user guide*.

# Amazon ECS endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ecs.us-east-2.amazonaws.com | HTTPS |
| | | ecs-fips.us-east-2.amazonaws.com | HTTPS |
| | | ecs-fips.us-east-2.api.aws | HTTPS |
| | | ecs.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | ecs.us-east-1.amazonaws.com | HTTPS |
| | | ecs-fips.us-east-1.amazonaws.com | HTTPS |
| | | ecs-fips.us-east-1.api.aws | HTTPS |
| | | ecs.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ecs.us-west-1.amazonaws.com | HTTPS |
| | | ecs-fips.us-west-1.amazonaws.com | HTTPS |
| | | ecs-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | ecs.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | ecs.us-west-2.amazonaws.com | HTTPS | |
| | | ecs-fips.us-west-2.amazonaws.com | HTTPS | |
| | | ecs-fips.us-west-2.api.aws | HTTPS | |
| | | ecs.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | ecs.af-south-1.amazonaws.com | HTTPS | |
| | | ecs.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | ecs.ap-east-1.amazonaws.com | HTTPS | |
| | | ecs.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | ecs.ap-south-2.amazonaws.com | HTTPS | |
| | | ecs.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ecs.ap-southeast-3.amazonaws.com | HTTPS | |
| | | ecs.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | ecs.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ecs.ap-southeast-4.amazonaws.com | HTTPS | |
| | | ecs.ap-southeast-4.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | ecs.ap-south-1.amazonaws.com<br>ecs.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ecs.ap-southeast-6.amazonaws.com<br>ecs.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ecs.ap-northeast-3.amazonaws.com<br>ecs.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ecs.ap-northeast-2.amazonaws.com<br>ecs.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | ecs.ap-southeast-1.amazonaws.com<br>ecs.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ecs.ap-southeast-2.amazonaws.com<br>ecs.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | ecs.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | ecs.ap-southeast-7.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | ecs.ap-northeast-1.amazonaws.com | HTTPS | |
| | | ecs.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | ecs.ca-central-1.amazonaws.com | HTTPS | |
| | | ecs.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | ecs.ca-west-1.amazonaws.com | HTTPS | |
| | | ecs.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | ecs.eu-central-1.amazonaws.com | HTTPS | |
| | | ecs.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | ecs.eu-west-1.amazonaws.com | HTTPS | |
| | | ecs.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | ecs.eu-west-2.amazonaws.com | HTTPS | |
| | | ecs.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | ecs.eu-south-1.amazonaws.com | HTTPS | |
| | | ecs.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | ecs.eu-west-3.amazonaws.com | HTTPS | |
| | | ecs.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | ecs.eu-south-2.amazonaws.com | HTTPS | |
| | | ecs.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | ecs.eu-north-1.amazonaws.com | HTTPS |
| | | ecs.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | ecs.eu-central-2.amazonaws.com | HTTPS |
| | | ecs.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | ecs.il-central-1.amazonaws.com | HTTPS |
| | | ecs.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | ecs.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | ecs.me-south-1.amazonaws.com | HTTPS |
| | | ecs.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | ecs.me-central-1.amazonaws.com | HTTPS |
| | | ecs.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | ecs.sa-east-1.amazonaws.com | HTTPS |
| | | ecs.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | ecs.us-gov-east-1.amazonaws.com | HTTPS |
| | | ecs-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | ecs-fips.us-gov-east-1.api.aws | HTTPS |
| | | ecs.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | ecs.us-gov-west-1.amazonaws.com | HTTPS | |
| | | ecs-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | ecs-fips.us-gov-west-1.api.aws | HTTPS | |
| | | ecs.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

The following are Amazon ECS service quotas.

Most of these service quotas, but not all, are listed under the Amazon Elastic Container Service (Amazon ECS) namespace in the Service Quotas console. To request a quota increase, see Requesting a quota increase in the *Service Quotas User Guide*.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Burst rate of agent modify actions (or bucket maximum capacity) | Each supported Region: 200 | Yes | The maximum burst number of requests per second for agent modify actions. |
| Burst rate of capacity provider modify actions (or bucket maximum capacity) | Each supported Region: 10 | Yes | The maximum burst number of requests per second for capacity provider modify actions. |
| Burst rate of capacity provider read actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for capacity provider read actions. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Burst rate of cluster modify actions (or bucket maximum capacity) | Each supported Region: 20 | Yes | The maximum burst number of requests per second for cluster modify actions. |
| Burst rate of cluster read actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for cluster read actions. |
| Burst rate of cluster resource modify actions (or bucket maximum capacity) | Each supported Region: 100 | Yes | The maximum burst number of requests per second for cluster resource modify actions. |
| Burst rate of cluster resource read actions (or bucket maximum capacity) | Each supported Region: 100 | Yes | The maximum burst number of requests per second for cluster resource read actions. |
| Burst rate of cluster service resource read actions (or bucket maximum capacity) | Each supported Region: 10 | Yes | The maximum burst number of requests per second for cluster service resource read actions. |
| Burst rate of service deployment read actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for service deployment and revision read actions. |
| Burst rate of service modify actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for service modify actions. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Burst rate of service read actions (or bucket maximum capacity) | Each supported Region: 100 | Yes | The maximum burst number of requests per second for service read actions. |
| Burst rate of setting modify actions (or bucket maximum capacity) | Each supported Region: 10 | Yes | The maximum burst number of requests per second for setting modify actions. |
| Burst rate of setting read actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for setting read actions. |
| Burst rate of tag modify actions (or bucket maximum capacity) | Each supported Region: 20 | Yes | The maximum burst number of requests per second for tag modify actions. |
| Burst rate of tag read actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for tag read actions. |
| Burst rate of task definition deletion actions (or bucket maximum capacity) | Each supported Region: 5 | Yes | The maximum burst number of requests per second for task definition deletion actions. |
| Burst rate of task definition modify actions (or bucket maximum capacity) | Each supported Region: 20 | Yes | The maximum burst number of requests per second for task definition modify actions. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Burst rate of task definition read actions (or bucket maximum capacity) | Each supported Region: 50 | Yes | The maximum burst number of requests per second for task definition read actions. |
| Burst rate of task protection actions (or bucket maximum capacity) | Each supported Region: 200 | Yes | The maximum burst number of requests per second for task protection write actions. |
| Burst rate of task protection actions (or bucket maximum capacity) | Each supported Region: 200 | Yes | The maximum burst number of requests per second for task protection read actions. |
| Capacity providers per cluster | Each supported Region: 20 | No | The maximum number of capacity providers that can be associated with a cluster. |
| Classic Load Balancers per service | Each supported Region: 1 | No | The maximum number of Classic Load Balancers per service. |
| Clusters per account | Each supported Region: 10,000 | Yes | Number of clusters per account |
| Container instances per cluster | Each supported Region: 5,000 | No | Number of container instances per cluster |
| Container instances per start-task | Each supported Region: 10 | No | The maximum number of container instances specified in a StartTask API action. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Containers per task definition | Each supported Region: 10 | No | The maximum number of containers definitions within a task definition. |
| ECS Exec sessions | Each supported Region: 1,000 | No | The maximum number of ECS Exec sessions per container. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of tasks launched by a service on AWS Fargate | us-east-1: 500<br>us-east-2: 500<br>us-west-1: 500<br>us-west-2: 500<br>ap-northeast-1: 500<br>ap-northeast-2: 500<br>ap-south-1: 500<br>ap-southeast-1: 500<br>ap-southeast-2: 500<br>ca-central-1: 500<br>eu-central-1: 500<br>eu-north-1: 500<br>eu-west-1: 500<br>eu-west-2: 500<br>eu-west-3: 500<br>sa-east-1: 500 | No | The maximum number of tasks that can be provisioned per service per minute on Fargate by the Amazon ECS service scheduler. Note: For Regions introduced after May 20, 2025, the maximum is 125 |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | Each of the other supported Regions: 125 | | |
| Rate of tasks launched by a service on an Amazon EC2 or External instance | Each supported Region: 500 | No | The maximum number of tasks that can be provisioned per service per minute on an Amazon EC2 or External instance by the Amazon ECS service scheduler. |
| Revisions per task definition family | Each supported Region: 1,000,000 | No | The maximum number of revisions per task definition family. Deregistering or deleting a task definition revision does not exclude it from being included in this limit. |
| Security groups per awsvpcConfiguration | Each supported Region: 5 | No | The maximum number of security groups specified within an awsvpcConfiguration. |
| Services per cluster | Each supported Region: 5,000 | No | The maximum number of services per cluster |
| Services per namespace | Each supported Region: 300 | No | The maximum number of services that can be running within a namespace. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Subnets per awsvpcConfiguration | Each supported Region: 16 | No | The maximum number of subnets specified within an awsvpcConfiguration. |
| Sustained rate of agent modify actions (or bucket refill rate) | Each supported Region: 120 | Yes | The maximum sustained number of requests per second for agent modify actions. |
| Sustained rate of capacity provider modify actions (or bucket refill rate) | Each supported Region: 1 | Yes | The maximum sustained number of requests per second for capacity provider modify actions. |
| Sustained rate of capacity provider read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for capacity provider read actions. |
| Sustained rate of cluster modify actions (or bucket refill rate) | Each supported Region: 1 | Yes | The maximum sustained number of requests per second for cluster modify actions. |
| Sustained rate of cluster read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for cluster read actions. |
| Sustained rate of cluster resource modify actions (or bucket refill rate) | Each supported Region: 40 | Yes | The maximum sustained number of requests per second for cluster resource modify actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sustained rate of cluster resource read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for cluster resource read actions. |
| Sustained rate of cluster service resource read actions (or bucket refill rate) | Each supported Region: 1 | Yes | The maximum sustained number of requests per second for cluster service resource read actions. |
| Sustained rate of service deployment read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for service deployment and revision read actions. |
| Sustained rate of service modify actions (or bucket refill rate) | Each supported Region: 5 | Yes | The maximum sustained number of requests per second for service modify actions. |
| Sustained rate of service read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for service read actions. |
| Sustained rate of setting modify actions (or bucket refill rate) | Each supported Region: 1 | Yes | The maximum sustained number of requests per second for setting modify actions. |
| Sustained rate of setting read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for setting read actions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sustained rate of tag modify actions (or bucket refill rate) | Each supported Region: 10 | Yes | The maximum sustained number of requests per second for tag modify actions. |
| Sustained rate of tag read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for tag read actions. |
| Sustained rate of task definition deletion actions (or bucket refill rate) | Each supported Region: 1 | Yes | The maximum sustained number of requests per second for task definition deletion actions. |
| Sustained rate of task definition modify actions (or bucket refill rate) | Each supported Region: 1 | Yes | The maximum sustained number of requests per second for task definition modify actions. |
| Sustained rate of task definition read actions (or bucket refill rate) | Each supported Region: 20 | Yes | The maximum sustained number of requests per second for task definition read actions. |
| Sustained rate of task protection actions (or bucket refill rate) | Each supported Region: 80 | Yes | The maximum sustained number of requests per second for task protection write actions. |
| Sustained rate of task protection actions (or bucket refill rate) | Each supported Region: 80 | Yes | The maximum sustained number of requests per second for task protection read actions. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Tags per resource | Each supported Region: 50 | No | The maximum number of tags per resource. This applies to task definitions, clusters, tasks, and services. |
| Target groups per service | Each supported Region: 5 | No | The maximum number of target groups per service, if using an Application Load Balancer or a Network Load Balancer. |
| Task definition size | Each supported Region: 64 Kilobytes | No | The maximum size, in KiB, of a task definition. |
| Tasks in PROVISIONING state per cluster | Each supported Region: 500 | No | The maximum number of tasks waiting in the PROVISIONING state per cluster. This quota only applies to tasks launched using an EC2 Auto Scaling group capacity provider. |
| Tasks launched per run-task | Each supported Region: 10 | No | The maximum number of tasks that can be launched per RunTask API action. |
| Tasks per service | Each supported Region: 5,000 | No | The maximum number of tasks per service (the desired count). |

**Note**

Services configured to use Amazon ECS service discovery have a limit of 1,000 tasks per service. This is due to the AWS Cloud Map service quota for the number of instances per service. For more information, see AWS Cloud Map service quotas in the *Amazon Web Services General Reference*.

**Note**

In practice, task launch rates are also dependent on other considerations such as container images to be downloaded and unpacked, health checks and other integrations enabled, such as registering tasks with a load balancer. You will see variations in task launch rates compared with the quotas represented above based on the features that you have enabled for your Amazon ECS services. For more information, see speeding up Amazon ECS deployments in the Amazon ECS Best Practices Guide.

## AWS Fargate quotas

The following are Amazon ECS on AWS Fargate service quotas and are listed under the **AWS Fargate** service in the Service Quotas console.

New AWS accounts might have initial lower quotas that can increase over time. Fargate constantly monitors the account usage within each Region, and then automatically increases the quotas based on your usage. You can also request a quota increase for values that are shown as adjustable, see Requesting a quota increase in the *Service Quotas User Guide*.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Fargate On-Demand Burst Launch Rate | us-east-1: 100<br>us-east-2: 100<br>us-west-1: 100<br>us-west-2: 100 | Yes | The maximum number of on-demand Amazon ECS tasks or Amazon EKS pods launched in a single burst (burst rate) in this account in the current |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | ap-northeast-1: 100<br><br>ap-northeast-2: 100<br><br>ap-south-1: 100<br><br>ap-southeast-1: 100<br><br>ap-southeast-2: 100<br><br>ca-central-1: 100<br><br>eu-central-1: 100<br><br>eu-north-1: 100<br><br>eu-west-1: 100<br><br>eu-west-2: 100<br><br>eu-west-3: 100<br><br>sa-east-1: 100<br><br>Each of the other supported Regions: 25 | | region. Note: For Regions introduced after May 20, 2025, the maximum is 25 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Fargate On-Demand Sustained Launch Rate | us-east-1: 20<br>us-east-2: 20<br>us-west-1: 20<br>us-west-2: 20<br>ap-northeast-1: 20<br>ap-northeast-2: 20<br>ap-south-1: 20<br>ap-southeast-1: 20<br>ap-southeast-2: 20<br>ca-central-1: 20<br>eu-central-1: 20<br>eu-north-1: 20<br>eu-west-1: 20<br>eu-west-2: 20<br>eu-west-3: 20<br>sa-east-1: 20 | Yes | The maximum number of on-demand Amazon ECS tasks or Amazon EKS pods launched per second (sustained rate) in this account in the current region. Note: For Regions introduced after May 20, 2025, the maximum is 5 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 5 | | |
| Fargate On-Demand vCPU resource count | Each supported Region: 6 | Yes | The number of Fargate vCPUs running concurrently as Fargate On-Demand in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Fargate Spot Burst Launch Rate | us-east-1: 100<br>us-east-2: 100<br>us-west-1: 100<br>us-west-2: 100<br>ap-northeast-1: 100<br>ap-northeast-2: 100<br>ap-south-1: 100<br>ap-southeast-1: 100<br>ap-southeast-2: 100<br>ca-central-1: 100<br>eu-central-1: 100<br>eu-north-1: 100<br>eu-west-1: 100<br>eu-west-2: 100<br>eu-west-3: 100<br>sa-east-1: 100 | Yes | The maximum number of Spot Amazon ECS tasks launched in a single burst (burst rate) in this account in the current region. Note: For Regions introduced after May 20, 2025, the maximum is 25 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 25 | | |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Fargate Spot Sustained Launch Rate | us-east-1: 20<br>us-east-2: 20<br>us-west-1: 20<br>us-west-2: 20<br>ap-northeast-1: 20<br>ap-northeast-2: 20<br>ap-south-1: 20<br>ap-southeast-1: 20<br>ap-southeast-2: 20<br>ca-central-1: 20<br>eu-central-1: 20<br>eu-north-1: 20<br>eu-west-1: 20<br>eu-west-2: 20<br>eu-west-3: 20<br>sa-east-1: 20 | Yes | The maximum number of Spot Amazon ECS tasks launched per second (sustained rate) in this account in the current region. Note: For Regions introduced after May 20, 2025, the maximum is 5 |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Each of the other supported Regions: 5 | | |
| Fargate Spot vCPU resource count | Each supported Region: 6 | Yes | The number of Fargate vCPUs running concurrently as Fargate Spot in this account in the current Region. |

**Note**

The default values are the initial quotas set by AWS, which are separate from the actual applied quota value and maximum possible service quota. For more information, see Terminology in Service Quotas in the *Service Quotas User Guide*.

**Note**

Fargate additionally enforces Amazon ECS tasks and Amazon EKS pods launch rate limits. For more information, see AWS Fargate throttling quotas in the *Amazon ECS Developer Guide*.

## Amazon Elastic Kubernetes Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | eks.us-east-2.amazonaws.com<br>fips.eks.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | eks.us-east-1.amazonaws.com<br>fips.eks.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | eks.us-west-1.amazonaws.com<br>fips.eks.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | eks.us-west-2.amazonaws.com<br>fips.eks.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | eks.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | eks.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | eks.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | eks.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | eks.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | eks.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | eks.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | eks.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | eks.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | eks.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | eks.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | eks.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | eks.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | eks.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | eks.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | eks.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | eks.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | eks.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | eks.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | eks.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | eks.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | eks.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | eks.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | eks.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | eks.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | eks.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | eks.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | eks.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | eks.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | eks.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | eks.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | eks.us-gov-west-1.amazonaws.com | HTTPS | |

Additionally, the EKS API has dual-stack endpoints. The dual-stack endpoint was introduced in August 2024. To use the dual-stack endpoints with the AWS CLI, see the Dual-stack and FIPS endpoints configuration in the *AWS SDKs and Tools Reference Guide*. The following lists the new endpoints:

EKS API public endpoint

```
eks.region.api.aws
```

IRSA OIDC Issuer URLs

```
oidc-eks.region.api.aws
```

The following endpoint provides the Amazon EKS Auth API. The EKS Pod Identity Agent on your nodes connects to this endpoint to get the credentials for each pod.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | eks-auth.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | eks-auth.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | eks-auth.us-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | eks-auth.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | eks-auth.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | eks-auth.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | eks-auth.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | eks-auth.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | eks-auth.ap-southeast-5.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | eks-auth.ap-southeast-4.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | eks-auth.ap-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | eks-auth.ap-southeast-6.api.aws | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | eks-auth.ap-northeast-3.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | eks-auth.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | eks-auth.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | eks-auth.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | eks-auth.ap-east-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | eks-auth.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | eks-auth.ap-northeast-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | eks-auth.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | eks-auth.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | eks-auth.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | eks-auth.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | eks-auth.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | eks-auth.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | eks-auth.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | eks-auth.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | eks-auth.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | eks-auth.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | eks-auth.il-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | eks-auth.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | eks-auth.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | eks-auth.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | eks-auth.sa-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | eks-auth.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | eks-auth.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Access entries per cluster | Each supported Region: 3,000 | No | The maximum number of access entries per cluster. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Clusters | Each supported Region: 100 | Yes | The maximum number of EKS clusters in this account in the current Region. |
| Control plane security groups per cluster | Each supported Region: 4 | No | The maximum number of control plane security groups per cluster (these are specified when you create the cluster). |
| EKS Anywhere Enterprise Subscriptions | Each supported Region: 10 | Yes | The maximum number of EKS Anywhere Enterprise Subscriptions in this account in the current Region. |
| Fargate profiles per cluster | Each supported Region: 10 | Yes | The maximum number of Fargate profiles per cluster. |
| Label pairs per Fargate profile selector | Each supported Region: 5 | Yes | The maximum number of label pairs per Fargate profile selector. |
| Managed node groups per cluster | Each supported Region: 30 | Yes | The maximum number of managed node groups per cluster. |
| Nodes per managed node group | Each supported Region: 450 | Yes | The maximum number of nodes per managed node group. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Public endpoint access CIDR ranges per cluster | Each supported Region: 40 | No | The maximum number of public endpoint access CIDR ranges per cluster (these are specified when you create or update the cluster). |
| Registered clusters | Each supported Region: 10 | Yes | The maximum number of registered clusters in this account in the current Region. |
| Remote node networks per cluster | Each supported Region: 15 | Yes | The maximum number of remote node networks per cluster. |
| Remote pod networks per cluster | Each supported Region: 15 | Yes | The maximum number of remote pod networks per cluster. |
| Selectors per Fargate profile | Each supported Region: 5 | Yes | The maximum number of selectors per Fargate profile. |

**Note**

The default values are the initial quotas set by AWS. These default values are separate from the actual applied quota values and maximum possible service quotas. For more information, see Terminology in Service Quotas in the *Service Quotas User Guide*.

These service quotas are listed under **Amazon Elastic Kubernetes Service (Amazon EKS)** in the Service Quotas console. To request a quota increase for values that are shown as adjustable, see Requesting a quota increase in the *Service Quotas User Guide*.

# AWS Fargate service quotas

The **AWS Fargate** service in the Service Quotas console lists several service quotas. The following table only describes the quota that is applicable to Amazon EKS.

New AWS accounts might have lower initial quotas that can increase over time. Fargate constantly monitors the account usage within each AWS Region, and then automatically increases the quotas based on the usage. You can also request a quota increase for values that are shown as adjustable. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Fargate On-Demand vCPU resource count | 6 | Yes | The number of Fargate vCPUs that can run concurrently as Fargate On-Demand in this account in the current Region. |

**Note**

The default values are the initial quotas set by AWS. These default values are separate from the actual applied quota values and maximum possible service quotas. For more information, see Terminology in Service Quotas in the *Service Quotas User Guide*.

**Note**

Fargate additionally enforces Amazon ECS tasks and Amazon EKS pods launch rate quotas. For more information, see AWS Fargate throttling quotas in the *Amazon Elastic Container Service User Guide for AWS Fargate*.

# Amazon Elastic File System endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | elasticfilesystem.us-east-2.amazonaws.com | HTTPS |
| | | efs.us-east-2.api.aws | HTTPS |
| | | efs-fips.us-east-2.api.aws | HTTPS |
| | | elasticfilesystem-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | elasticfilesystem.us-east-1.amazonaws.com | HTTPS |
| | | efs.us-east-1.api.aws | HTTPS |
| | | efs-fips.us-east-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | elasticfilesystem.us-west-1.amazonaws.com | HTTPS |
| | | efs.us-west-1.api.aws | HTTPS |
| | | efs-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | elasticfilesystem-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | elasticfilesystem.us-west-2.amazonaws.com | HTTPS | |
| | | efs.us-west-2.api.aws | HTTPS | |
| | | efs-fips.us-west-2.api.aws | HTTPS | |
| | | elasticfilesystem-fips.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | elasticfilesystem.af-south-1.amazonaws.com | HTTPS | |
| | | efs.af-south-1.api.aws | HTTPS | |
| | | efs-fips.af-south-1.api.aws | HTTPS | |
| | | elasticfilesystem-fips.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | elasticfilesystem.ap-east-1.amazonaws.com | HTTPS | |
| | | efs.ap-east-1.api.aws | HTTPS | |
| | | efs-fips.ap-east-1.api.aws | HTTPS | |
| | | elasticfilesystem-fips.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | elasticfilesystem.ap-south-2.amazonaws.com | HTTPS | |
| | | efs.ap-south-2.api.aws | HTTPS | |
| | | efs-fips.ap-south-2.api.aws | HTTPS | |
| | | elasticfilesystem-fips.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | elasticfilesystem.ap-southeast-3.amazonaws.com<br>efs.ap-southeast-3.api.aws<br>efs-fips.ap-southeast-3.api.aws<br>elasticfilesystem-fips.ap-southeast-3.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | elasticfilesystem.ap-southeast-5.amazonaws.com<br>efs.ap-southeast-5.api.aws<br>efs-fips.ap-southeast-5.api.aws<br>elasticfilesystem-fips.ap-southeast-5.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | elasticfilesystem.ap-southeast-4.amazonaws.com<br>efs.ap-southeast-4.api.aws<br>efs-fips.ap-southeast-4.api.aws<br>elasticfilesystem-fips.ap-southeast-4.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | elasticfilesystem.ap-south-1.amazonaws.com<br>efs.ap-south-1.api.aws<br>efs-fips.ap-south-1.api.aws<br>elasticfilesystem-fips.ap-south-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | elasticfilesystem.ap-southeast-6.amazonaws.com<br>efs.ap-southeast-6.api.aws<br>efs-fips.ap-southeast-6.api.aws<br>elasticfilesystem-fips.ap-southeast-6.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | elasticfilesystem.ap-northeast-3.amazonaws.com<br>efs.ap-northeast-3.api.aws<br>efs-fips.ap-northeast-3.api.aws<br>elasticfilesystem-fips.ap-northeast-3.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | elasticfilesystem.ap-northeast-2.amazonaws.com<br>efs.ap-northeast-2.api.aws<br>efs-fips.ap-northeast-2.api.aws<br>elasticfilesystem-fips.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | elasticfilesystem.ap-southeast-1.amazonaws.com<br>efs.ap-southeast-1.api.aws<br>efs-fips.ap-southeast-1.api.aws<br>elasticfilesystem-fips.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | elasticfilesystem.ap-southeast-2.amazonaws.com<br>efs.ap-southeast-2.api.aws<br>efs-fips.ap-southeast-2.api.aws<br>elasticfilesystem-fips.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | elasticfilesystem.ap-east-2.amazonaws.com<br>elasticfilesystem-fips.ap-east-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | elasticfilesystem.ap-southeast-7.amazonaws.com<br>efs.ap-southeast-7.api.aws<br>efs-fips.ap-southeast-7.api.aws<br>elasticfilesystem-fips.ap-southeast-7.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | elasticfilesystem.ap-northeast-1.amazonaws.com | HTTPS |
| | | efs.ap-northeast-1.api.aws | HTTPS |
| | | efs-fips.ap-northeast-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | elasticfilesystem.ca-central-1.amazonaws.com | HTTPS |
| | | efs.ca-central-1.api.aws | HTTPS |
| | | efs-fips.ca-central-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | elasticfilesystem.ca-west-1.amazonaws.com | HTTPS |
| | | efs.ca-west-1.api.aws | HTTPS |
| | | efs-fips.ca-west-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.ca-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | elasticfilesystem.eu-central-1.amazonaws.com | HTTPS |
| | | efs.eu-central-1.api.aws | HTTPS |
| | | efs-fips.eu-central-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.eu-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | elasticfilesystem.eu-west-1.amazonaws.com | HTTPS | |
| | | efs.eu-west-1.api.aws | HTTPS | |
| | | efs-fips.eu-west-1.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | elasticfilesystem.eu-west-2.amazonaws.com | HTTPS | |
| | | efs.eu-west-2.api.aws | HTTPS | |
| | | efs-fips.eu-west-2.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | elasticfilesystem.eu-south-1.amazonaws.com | HTTPS | |
| | | efs.eu-south-1.api.aws | HTTPS | |
| | | efs-fips.eu-south-1.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | elasticfilesystem.eu-west-3.amazonaws.com | HTTPS | |
| | | efs.eu-west-3.api.aws | HTTPS | |
| | | efs-fips.eu-west-3.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | elasticfilesystem.eu-south-2.amazonaws.com | HTTPS | |
| | | efs.eu-south-2.api.aws | HTTPS | |
| | | efs-fips.eu-south-2.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | elasticfilesystem.eu-north-1.amazonaws.com | HTTPS | |
| | | efs.eu-north-1.api.aws | HTTPS | |
| | | efs-fips.eu-north-1.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | elasticfilesystem.eu-central-2.amazonaws.com | HTTPS | |
| | | efs.eu-central-2.api.aws | HTTPS | |
| | | efs-fips.eu-central-2.api.aws | HTTPS | |
| | | elasticfilesystem-fips.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | elasticfilesystem.il-central-1.amazonaws.com | HTTPS | |
| | | efs.il-central-1.api.aws | HTTPS | |
| | | efs-fips.il-central-1.api.aws | HTTPS | |
| | | elasticfilesystem-fips.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | elasticfilesystem.mx-central-1.amazonaws.com | HTTPS |
| | | efs.mx-central-1.api.aws | HTTPS |
| | | efs-fips.mx-central-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | elasticfilesystem.me-south-1.amazonaws.com | HTTPS |
| | | efs.me-south-1.api.aws | HTTPS |
| | | efs-fips.me-south-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | elasticfilesystem.me-central-1.amazonaws.com | HTTPS |
| | | efs.me-central-1.api.aws | HTTPS |
| | | efs-fips.me-central-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | elasticfilesystem.sa-east-1.amazonaws.com | HTTPS |
| | | efs.sa-east-1.api.aws | HTTPS |
| | | efs-fips.sa-east-1.api.aws | HTTPS |
| | | elasticfilesystem-fips.sa-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | elasticfilesystem.us-gov-east-1.amazonaws.com<br>efs.us-gov-east-1.api.aws<br>efs-fips.us-gov-east-1.api.aws<br>elasticfilesystem-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | elasticfilesystem.us-gov-west-1.amazonaws.com<br>efs.us-gov-west-1.api.aws<br>efs-fips.us-gov-west-1.api.aws<br>elasticfilesystem-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Access points per file system | Each supported Region: 10,000 | No | The maximum number of access points allowed per file system |
| Active users per NFS client | Each supported Region: 65,536 | No | The maximum number of active users that can have files open at the same time for each NFS client |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Bursting throughput | us-east-1: 3,072 Megabytes per second<br><br>us-east-2: 3,072 Megabytes per second<br><br>us-west-2: 3,072 Megabytes per second<br><br>ap-southeast-2: 3,072 Megabytes per second<br><br>eu-west-1: 3,072 Megabytes per second<br><br>Each of the other supported Regions: 1,024 Megabytes per second | No | The maximum total bursting throughput for all connected clients on an EFS file system |
| Directory depth | Each supported Region: 1,000 | No | The maximum number of levels, or depth, that you can have within a single directory |
| EFS file locks | Each supported Region: 512 | No | The maximum number of locks per file across all instances connected and users accessing the file |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| File hard links | Each supported Region: 177 | No | The maximum number of hard links that you can have for each file |
| File size | Each supported Region: 52,673,613,135,872 Bytes | No | The maximum size of a single file on an EFS file system |
| File system name length | Each supported Region: 255 Bytes | No | The maximum length of a file system name |
| File system symbolic link (symlink) length | Each supported Region: 4,080 Bytes | No | The maximum length that a file system symbolic link, or symlink, can be |
| File systems per account | Each supported Region: 1,000 | Yes | The maximum number of file systems that a customer account can have in an AWS Region |
| Locks across unique file/process pairs | Each supported Region: 65,536 | No | The maximum number of locks that can occur at the same time for each unique mount on an NFS client, across all unique file/process pairs |
| Maximum read IOPS per file system with Elastic throughput | Each supported Region: 250,000 | Yes | The maximum read IOPS file systems using Elastic throughput can achieve. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum wait time between Provisioned Throughput decreases | Each supported Region: 86,400 Seconds | No | The minimum amount of time that you have to wait after decreasing the amount of Provisioned Throughput before you can decrease it again |
| Minimum wait time between Throughput mode changes | Each supported Region: 86,400 Seconds | No | The minimum amount of time that you have to wait after changing the Throughput mode before you can change it again |
| Mount targets per Availability Zone | Each supported Region: 1 | No | The maximum number of EFS mount targets that you can have in an Availability Zone |
| Mount targets per VPC | Each supported Region: 1,400 | No | The maximum number of EFS mount targets allowed for each VPC |
| Open files per NFS client | Each supported Region: 65,536 | No | The maximum number of files that can be open at the same time for each NFS client |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Provisioned throughput | us-east-1: 3,414 Megabytes per second<br><br>us-east-2: 3,414 Megabytes per second<br><br>us-west-2: 3,414 Megabytes per second<br><br>eu-west-1: 3,414 Megabytes per second<br><br>Each of the other supported Regions: 1,024 Megabytes per second | No | The maximum total provisioned throughput for all connected clients on an EFS file system |
| Rate of file system operations | Each supported Region: 250,000 | No | The maximum number of file system operations per second, for all connected clients, on an EFS file system that is configured in General Purpose mode. Data or metadata read operations consume one file operation, and data or metadata write operations consume five file operations. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Security groups per mount target | Each supported Region: 5 | No | The maximum number of security groups that you can apply to a single EFS mount target |
| Tags | Each supported Region: 50 | No | The maximum number of tags that you can apply to an EFS file system |
| Throughput per NFS client | Each supported Region: 1,500 Megabytes per second | No | The maximum throughput that you can drive for each NFS client connected to an EFS file system |
| VPCs per file system | Each supported Region: 1 | No | The maximum number of VPCs that you can apply to an EFS file system |

For more information, see Amazon EFS quotas in the *Amazon Elastic File System User Guide*.

## Elastic Load Balancing endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Route 53 Hosted Zone ID (Application Load Balancers, Classic Load Balancers) | Route 53 Hosted Zone ID (Network Load Balancers) |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | elasticloadbalancing.us-east-2.amazonaws.com<br><br>elasticloadbalancing-fips.us-east-2.amazonaws.com | Z3AADJGX6KTTL2 | ZLMOA37VPKANP |
| US East (N. Virginia) | us-east-1 | elasticloadbalancing.us-east-1.amazonaws.com<br><br>elasticloadbalancing-fips.us-east-1.amazonaws.com | Z35SXDOTRQ7X7K | Z26RNL4JYFTOTI |
| US West (N. California) | us-west-1 | elasticloadbalancing.us-west-1.amazonaws.com<br><br>elasticloadbalancing-fips.us-west-1.amazonaws.com | Z368ELLRRE2KJ0 | Z24FKFUX50B4VW |
| US West (Oregon) | us-west-2 | elasticloadbalancing.us-west-2.amazonaws.com<br><br>elasticloadbalancing-fips.us-west-2.amazonaws.com | Z1H1FL5HABSF5 | Z18D5FSROUN65G |
| Africa (Cape Town) | af-south-1 | elasticloadbalancing.af-south-1.amazonaws.com | Z268VQBMOI5EKX | Z203XCE67M25HM |

| Region Name | Region | Endpoint | Route 53 Hosted Zone ID (Application Load Balancers, Classic Load Balancers) | Route 53 Hosted Zone ID (Network Load Balancers) |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | elasticloadbalancing.ap-east-1.amazonaws.com | Z3DQVH9N71FHZ0 | Z12Y7K3UBGUAD1 |
| Asia Pacific (Hyderabad) | ap-south-2 | elasticloadbalancing.ap-south-2.amazonaws.com | Z0173938T07WNTVAEPZN | Z0711778386UTO08407HT |
| Asia Pacific (Jakarta) | ap-southeast-3 | elasticloadbalancing.ap-southeast-3.amazonaws.com | Z08888821HLRG5A9ZRTER | Z01971771FYVNCOVWJU1G |
| Asia Pacific (Malaysia) | ap-southeast-5 | elasticloadbalancing.ap-southeast-5.amazonaws.com | Z06010284QMVVW7WO5J | Z026317210H9ACVTRO6FB |
| Asia Pacific (Melbourne) | ap-southeast-4 | elasticloadbalancing.ap-southeast-4.amazonaws.com | Z09517862IB2WZLPXG76F | Z01156963G8MIIL7X90IV |
| Asia Pacific (Mumbai) | ap-south-1 | elasticloadbalancing.ap-south-1.amazonaws.com | ZP97RAFLXTNZK | ZVDDRBQ08TROA |
| Asia Pacific (New Zealand) | ap-southeast-6 | elasticloadbalancing.ap-southeast-6.amazonaws.com | Z023301818UFJ50CIO0MV | Z01392953RKV2Q3RBP0KU |
| Asia Pacific (Osaka) | ap-northeast-3 | elasticloadbalancing.ap-northeast-3.amazonaws.com | Z5LXEXXYW11ES | Z1GWIQ4HH19I5X |

| Region Name | Region | Endpoint | Route 53 Hosted Zone ID (Application Load Balancers, Classic Load Balancers) | Route 53 Hosted Zone ID (Network Load Balancers) |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | elasticloadbalancing.ap-northeast-2.amazonaws.com | ZWKZPGTI48KDX | ZIBE1TIR4HY56 |
| Asia Pacific (Singapore) | ap-southeast-1 | elasticloadbalancing.ap-southeast-1.amazonaws.com | Z1LMS91P8CMLE5 | ZKVM4W9LS7TM |
| Asia Pacific (Sydney) | ap-southeast-2 | elasticloadbalancing.ap-southeast-2.amazonaws.com | Z1GM3OXH4ZPM65 | ZCT6FZBF4DROD |
| Asia Pacific (Taipei) | ap-east-2 | elasticloadbalancing.ap-east-2.amazonaws.com | Z02789141MW7T1WBU19PO | Z09176273OC2HWIAUNYW |
| Asia Pacific (Thailand) | ap-southeast-7 | elasticloadbalancing.ap-southeast-7.amazonaws.com | Z0390008CMBRTHFGWBCB | Z054363131YWATEMWRG5L |
| Asia Pacific (Tokyo) | ap-northeast-1 | elasticloadbalancing.ap-northeast-1.amazonaws.com | Z14GRHDCWA56QT | Z31USIVHYNEOWT |
| Canada (Central) | ca-central-1 | elasticloadbalancing.ca-central-1.amazonaws.com | ZQSVJUPU6J1EY | Z2EPGBW3API2WT |
| Canada West (Calgary) | ca-west-1 | elasticloadbalancing.ca-west-1.amazonaws.com | Z06473681N0SF6OS049SD | Z02754302KBB00W2LKWZ9 |

| Region Name | Region | Endpoint | Route 53 Hosted Zone ID (Application Load Balancers, Classic Load Balancers) | Route 53 Hosted Zone ID (Network Load Balancers) |
|---|---|---|---|---|
| China (Beijing) | cn-north-1 | elasticloadbalancing.cn-north-1.amazonaws.com.cn | Z1GDH35T77C1KE | Z3QFB96KMJ7ED6 |
| China (Ningxia) | cn-northwest-1 | elasticloadbalancing.cn-northwest-1.amazonaws.com.cn | ZM7IZAIOVVDZF | ZQEIKTCZ8352D |
| Europe (Frankfurt) | eu-central-1 | elasticloadbalancing.eu-central-1.amazonaws.com | Z215JYRZR1TBD5 | Z3F0SRJ5LGBH90 |
| Europe (Ireland) | eu-west-1 | elasticloadbalancing.eu-west-1.amazonaws.com | Z32O12XQLNTSW2 | Z2IFOLAFXWLO4F |
| Europe (London) | eu-west-2 | elasticloadbalancing.eu-west-2.amazonaws.com | ZHURV8PSTC4K8 | ZD4D7Y8KGAS4G |
| Europe (Milan) | eu-south-1 | elasticloadbalancing.eu-south-1.amazonaws.com | Z3ULH7SSC9OV64 | Z23146JA1KNAFP |
| Europe (Paris) | eu-west-3 | elasticloadbalancing.eu-west-3.amazonaws.com | Z3Q77PNBQS71R4 | Z1CMS0P5QUZ6D5 |
| Europe (Spain) | eu-south-2 | elasticloadbalancing.eu-south-2.amazonaws.com | Z0956581394HF5D5LXGAP | Z1011216NVTVYADP1SSV |
| Europe (Stockholm) | eu-north-1 | elasticloadbalancing.eu-north-1.amazonaws.com | Z23TAZ6LKFMNIO | Z1UDT6IFJ4EJM |

| Region Name | Region | Endpoint | Route 53 Hosted Zone ID (Application Load Balancers, Classic Load Balancers) | Route 53 Hosted Zone ID (Network Load Balancers) |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | elasticloadbalancing.eu-central-2.amazonaws.com | Z06391101F2ZOEP8P5EB3 | Z02239872DOALSIDCX66S |
| Israel (Tel Aviv) | il-central-1 | elasticloadbalancing.il-central-1.amazonaws.com | Z09170902867EHPV2DABU | Z0313266YDI6ZRHTGQY4 |
| Mexico (Central) | mx-central-1 | elasticloadbalancing.mx-central-1.amazonaws.com | Z023552324OKD1BB28BH5 | Z02031231H3ID6HYJ9A7U |
| Middle East (Bahrain) | me-south-1 | elasticloadbalancing.me-south-1.amazonaws.com | ZS929ML54UICD | Z3QSRYVP46NYYV |
| Middle East (UAE) | me-central-1 | elasticloadbalancing.me-central-1.amazonaws.com | Z08230872XQRWHG2XF6I | Z00282643NTTLPANJJG2P |
| South America (São Paulo) | sa-east-1 | elasticloadbalancing.sa-east-1.amazonaws.com | Z2P70J7HTTTPLU | ZTK26PT1VY4CU |
| AWS GovCloud (US-East) | us-gov-east-1 | elasticloadbalancing.us-gov-east-1.amazonaws.com | Z166TLBEWOO7G0 | Z1ZSMQQ6Q24QQ8 |
| AWS GovCloud (US-West) | us-gov-west-1 | elasticloadbalancing.us-gov-west-1.amazonaws.com | Z33AYJ8TM3BH4J | ZMG1MZ2THAWF1 |

# Service quotas

The following quotas are for Application Load Balancers.

| Name | Default | Adjustable |
|---|---|---|
| Application Load Balancers per Region | 50 | Yes |
| Certificates per Application Load Balancer | 25 | Yes |
| Condition Values per Rule | 5 | No |
| Condition Wildcards per Rule | 6 | No |
| Listeners per Application Load Balancer | 50 | Yes |
| Reserved Application Load Balancer Capacity Units (LCU) per Region | 0 | Yes |
| Rules per Application Load Balancer | 100 | Yes |
| Target Groups per Action per Application Load Balancer | 5 | No |
| Target Groups per Application Load Balancer | 100 | No |
| Targets per Application Load Balancer | 1,000 | Yes |

The following quotas are for Network Load Balancers.

| Name | Default | Adjustable |
|---|---|---|
| Certificates per Network Load Balancer | 25 | Yes |
| Listeners per Network Load Balancer | 50 | No |
| Network Load Balancer ENIs per VPC | 1,200 | Yes |
| Network Load Balancers per Region | 50 | Yes |

| Name | Default | Adjustable |
| --- | --- | --- |
| Reserved Network Load Balancer Capacity Units (LCU) per Region | 0 | Yes |
| Targets per Availability Zone per Network Load Balancer | 500 | Yes |
| Targets per Network Load Balancer | 3,000 | Yes |

The following quotas are for Gateway Load Balancers.

| Name | Default | Adjustable |
| --- | --- | --- |
| Gateway Load Balancers per Region | 100 | Yes |
| Gateway Load Balancers per VPC | 100 | Yes |
| Gateway Load Balancer ENIs per VPC | 300 * | Yes |
| GENEVE target groups per Region | 100 | Yes |
| Listeners per Gateway Load Balancer | 1 | No |
| Reserved Gateway Load Balancer Capacity Units (LCU) per Region | 0 | Yes |
| Targets per Availability Zone per GENEVE target group | 300 | No |
| Targets per Availability Zone per Gateway Load Balancer | 300 | No |
| Targets per Gateway Load Balancer | 300 | No |

The following quotas are for target groups.

| Name | Default | Adjustable |
| --- | --- | --- |
| Target Groups per Region | 3,000 | Yes |
| Targets per Target Group per Region | 1,000 | Yes |

The following quotas are for Classic Load Balancers.

| Name | Default | Adjustable |
| --- | --- | --- |
| Classic Load Balancers per Region | 20 | Yes |
| Listeners per Classic Load Balancer | 100 | Yes |
| Registered Instances per Classic Load Balancer | 1,000 | Yes |

For more information, see the following:

- Quotas for your Application Load Balancers
- Quotas for your Network Load Balancers
- Quotas for your Classic Load Balancers
- Quotas for your Gateway Load Balancers

## AWS Elastic Disaster Recovery endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | drs.us-east-2.amazonaws.com | HTTPS | |
| | | drs-fips.us-east-2.amazonaws.com | HTTPS | |
| | | drs-fips.us-east-2.api.aws | HTTPS | |
| | | drs.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | drs.us-east-1.amazonaws.com | HTTPS | |
| | | drs-fips.us-east-1.amazonaws.com | HTTPS | |
| | | drs-fips.us-east-1.api.aws | HTTPS | |
| | | drs.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | drs.us-west-1.amazonaws.com | HTTPS | |
| | | drs-fips.us-west-1.amazonaws.com | HTTPS | |
| | | drs-fips.us-west-1.api.aws | HTTPS | |
| | | drs.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | drs.us-west-2.amazonaws.com | HTTPS | |
| | | drs-fips.us-west-2.amazonaws.com | HTTPS | |
| | | drs-fips.us-west-2.api.aws | HTTPS | |
| | | drs.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | drs.af-south-1.amazonaws.com | HTTPS | |
| | | drs.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | drs.ap-east-1.amazonaws.com<br>drs.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | drs.ap-south-2.amazonaws.com<br>drs.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | drs.ap-southeast-3.amazonaws.com<br>drs.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | drs.ap-southeast-5.amazonaws.com<br>drs.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | drs.ap-southeast-4.amazonaws.com<br>drs.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | drs.ap-south-1.amazonaws.com<br>drs.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | drs.ap-southeast-6.amazonaws.com<br>drs.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | drs.ap-northeast-3.amazonaws.com<br>drs.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | drs.ap-northeast-2.amazonaws.com<br>drs.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | drs.ap-southeast-1.amazonaws.com<br>drs.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | drs.ap-southeast-2.amazonaws.com<br>drs.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | drs.ap-east-2.amazonaws.com<br>drs.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | drs.ap-southeast-7.amazonaws.com<br>drs.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | drs.ap-northeast-1.amazonaws.com<br>drs.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | drs.ca-central-1.amazonaws.com<br>drs.ca-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | drs.ca-west-1.amazonaws.com | HTTPS | |
| | | drs.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | drs.eu-central-1.amazonaws.com | HTTPS | |
| | | drs.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | drs.eu-west-1.amazonaws.com | HTTPS | |
| | | drs.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | drs.eu-west-2.amazonaws.com | HTTPS | |
| | | drs.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | drs.eu-south-1.amazonaws.com | HTTPS | |
| | | drs.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | drs.eu-west-3.amazonaws.com | HTTPS | |
| | | drs.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | drs.eu-south-2.amazonaws.com | HTTPS | |
| | | drs.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | drs.eu-north-1.amazonaws.com | HTTPS | |
| | | drs.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | drs.eu-central-2.amazonaws.com | HTTPS | |
| | | drs.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | drs.il-central-1.amazonaws.com | HTTPS |
| | | drs.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | drs.mx-central-1.amazonaws.com | HTTPS |
| | | drs.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | drs.me-south-1.amazonaws.com | HTTPS |
| | | drs.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | drs.me-central-1.amazonaws.com | HTTPS |
| | | drs.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | drs.sa-east-1.amazonaws.com | HTTPS |
| | | drs.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | drs.us-gov-east-1.amazonaws.com | HTTPS |
| | | drs-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | drs-fips.us-gov-east-1.api.aws | HTTPS |
| | | drs.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | drs.us-gov-west-1.amazonaws.com | HTTPS |
| | | drs-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | drs-fips.us-gov-west-1.api.aws | HTTPS |
| | | drs.us-gov-west-1.api.aws | HTTPS |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Concurrent jobs in progress | Each supported Region: 20 | No | Concurrent jobs in progress |
| Max Total replicating source servers Per AWS Account | Each supported Region: 300 | No | Max Total replicating source servers Per AWS Account |
| Max Total source servers Per AWS Account | Each supported Region: 4,000 | Yes | Max Total source servers Per AWS Account |
| Max concurrent Jobs per source server | Each supported Region: 1 | No | Max concurrent Jobs per source server |
| Max number of launch actions per resource | Each supported Region: 200 | No | Max number of launch actions per resource |
| Max number of launch configuration templates per AWS account | Each supported Region: 1 | No | Max number of launch configuration templates per AWS account |
| Max number of source networks per AWS account | Each supported Region: 100 | Yes | Max number of source networks per AWS account |
| Max source servers in a single Job | Each supported Region: 100 | No | Max source servers in a single Job |
| Max source servers in all Jobs | Each supported Region: 500 | No | Max source servers in all Jobs |

# Amazon ElastiCache endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | elasticache.us-east-2.amazonaws.com | HTTPS |
| | | elasticache-fips.us-east-2.api.aws | HTTPS |
| | | elasticache.us-east-2.api.aws | HTTPS |
| | | elasticache-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | elasticache.us-east-1.amazonaws.com | HTTPS |
| | | elasticache-fips.us-east-1.api.aws | HTTPS |
| | | elasticache-fips.us-east-1.amazonaws.com | HTTPS |
| | | elasticache.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | elasticache.us-west-1.amazonaws.com | HTTPS |
| | | elasticache.us-west-1.api.aws | HTTPS |
| | | elasticache-fips.us-west-1.amazonaws.com | HTTPS |
| | | elasticache-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | elasticache.us-west-2.amazonaws.com | HTTPS |
| | | elasticache-fips.us-west-2.api.aws | HTTPS |
| | | elasticache-fips.us-west-2.amazonaws.com | HTTPS |
| | | elasticache.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | elasticache.af-south-1.amazonaws.com | HTTPS |
| | | elasticache.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | elasticache.ap-east-1.amazonaws.com | HTTPS |
| | | elasticache.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | elasticache.ap-south-2.amazonaws.com | HTTPS |
| | | elasticache.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | elasticache.ap-southeast-3.amazonaws.com | HTTPS |
| | | elasticache.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | elasticache.ap-southeast-5.amazonaws.com | HTTPS |
| | | elasticache.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | elasticache.ap-southeast-4.amazonaws.com | HTTPS |
| | | elasticache.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | elasticache.ap-south-1.amazonaws.com<br>elasticache.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | elasticache.ap-southeast-6.amazonaws.com<br>elasticache.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | elasticache.ap-northeast-3.amazonaws.com<br>elasticache.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | elasticache.ap-northeast-2.amazonaws.com<br>elasticache.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | elasticache.ap-southeast-1.amazonaws.com<br>elasticache.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | elasticache.ap-southeast-2.amazonaws.com<br>elasticache.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | elasticache.ap-east-2.amazonaws.com<br>elasticache.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | elasticache.ap-southeast-7.amazonaws.com<br>elasticache.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | elasticache.ap-northeast-1.amazonaws.com | HTTPS |
| | | elasticache.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | elasticache.ca-central-1.amazonaws.com | HTTPS |
| | | elasticache.ca-central-1.api.aws | HTTPS |
| | | elasticache-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | elasticache.ca-west-1.amazonaws.com | HTTPS |
| | | elasticache.ca-west-1.api.aws | HTTPS |
| | | elasticache-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | elasticache.eu-central-1.amazonaws.com | HTTPS |
| | | elasticache.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | elasticache.eu-west-1.amazonaws.com | HTTPS |
| | | elasticache.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | elasticache.eu-west-2.amazonaws.com | HTTPS |
| | | elasticache.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | elasticache.eu-south-1.amazonaws.com | HTTPS |
| | | elasticache.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | elasticache.eu-west-3.amazonaws.com | HTTPS |
| | | elasticache.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | elasticache.eu-south-2.amazonaws.com | HTTPS |
| | | elasticache.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | elasticache.eu-north-1.amazonaws.com | HTTPS |
| | | elasticache.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | elasticache.eu-central-2.amazonaws.com | HTTPS |
| | | elasticache.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | elasticache.il-central-1.amazonaws.com | HTTPS |
| | | elasticache.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | elasticache.mx-central-1.amazonaws.com | HTTPS |
| | | elasticache.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | elasticache.me-south-1.amazonaws.com | HTTPS |
| | | elasticache.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | elasticache.me-central-1.amazonaws.com | HTTPS |
| | | elasticache.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | elasticache.sa-east-1.amazonaws.com | HTTPS |
| | | elasticache.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | elasticache.us-gov-east-1.amazonaws.com | HTTPS |
| | | elasticache.us-gov-east-1.api.aws | HTTPS |
| | | elasticache-fips.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | elasticache.us-gov-west-1.amazonaws.com<br>elasticache.us-gov-west-1.api.aws<br>elasticache-fips.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Nodes per Region | Each supported Region: 300 | Yes | The maximum number of nodes across all clusters in a Region. This quota applies to both your reserved and non-reserved nodes within the given Region. You can have up to 300 reserved nodes and 300 non-reserved nodes in the same Region. |
| Nodes per cluster (Memcached) | Each supported Region: 60 | Yes | The maximum number of nodes in an individual Memcached cluster. |
| Nodes per cluster (cluster mode enabled) | Each supported Region: 90 | Yes | The maximum number of nodes in an individual Redis or Valkey cluster. |
| Nodes per shard | Each supported Region: 6 | No | The maximum number of nodes in an individual shard (node group). One |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | node is the read/write Primary. All other nodes are read-only Replicas. |
| Parameter groups per Region | Each supported Region: 300 | Yes | The maximum number of parameters groups you can create in a Region. |
| Serverless Caches per Region | Each supported Region: 40 | Yes | The maximum number of serverless caches in a Region. |
| Serverless snapshots per day per cache | Each supported Region: 24 | Yes | The maximum number of snapshots that can be taken per day per serverless cache. |
| Subnet groups per Region | Each supported Region: 300 | Yes | The maximum number of subnet groups you can create in a Region. |
| Subnets per subnet group | Each supported Region: 20 | Yes | The maximum number of subnets you can define for a subnet group. |
| User Groups per Region | Each supported Region: 200 | Yes | The maximum number of user groups you can create in a Region. |
| Users per Region | Each supported Region: 2,000 | Yes | The maximum number of users you can create in a Region. |
| Users per User Group | Each supported Region: 100 | Yes | The maximum number of users you can define for a user group. |

# AWS Elemental Inference endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | elemental-inference.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | elemental-inference.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | elemental-inference.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | elemental-inference.ap-southeast-2.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | elemental-inference.eu-west-1.amazonaws.com | HTTPS | |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Active feeds per account | ap-south-1: 1 Count<br><br>eu-west-1: 1 Count<br><br>Each of the other supported Regions: 2 Count | Yes | The maximum number of feeds that can be actively running simultaneously in this account in the current AWS Region. |
| Enabled outputs per feed | Each supported Region: 3 Count | No | The maximum number of outputs that can be enabled per feed in this account in the current AWS Region. |
| Request rate for AssociateFeed | Each supported Region: 5 Count/ Second | No | The maximum number of AssociateFeed requests per second that you can send in this account in the current AWS Region. |
| Request rate for AssociateFeed, in a burst | Each supported Region: 30 Count | No | The maximum number of AssociateFeed requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for CreateFeed | Each supported Region: 5 Count/ Second | No | The maximum number of CreateFeed requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for CreateFeed, in a burst | Each supported Region: 30 Count | No | The maximum number of CreateFeed requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for DeleteFeed | Each supported Region: 5 Count/ Second | No | The maximum number of DeleteFeed requests per second that you can send in this account in the current AWS Region. |
| Request rate for DeleteFeed, in a burst | Each supported Region: 30 Count | No | The maximum number of DeleteFeed requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for DisassociateFeed | Each supported Region: 5 Count/ Second | No | The maximum number of DisassociateFeed requests per second that you can send in this account in the current AWS Region. |
| Request rate for DisassociateFeed, in a burst | Each supported Region: 30 Count | No | The maximum number of DisassociateFeed requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for GetFeed | Each supported Region: 5 Count/Second | No | The maximum number of GetFeed requests per second that you can send in this account in the current AWS Region. |
| Request rate for GetFeed, in a burst | Each supported Region: 30 Count | No | The maximum number of GetFeed requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for GetMetadata | Each supported Region: 2 Count/Second | No | The maximum number of GetMetadata requests per second that you can send for each feed, input, and output combination in the current AWS Region. |
| Request rate for GetMetadata, in a burst | Each supported Region: 5 Count | No | The maximum number of GetMetadata requests that you can send in one burst for each feed, input, and output combination in the current AWS Region. |
| Request rate for ListFeeds | Each supported Region: 5 Count/Second | No | The maximum number of ListFeeds requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for ListFeeds, in a burst | Each supported Region: 30 Count | No | The maximum number of ListFeeds requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for PutMedia | Each supported Region: 1.5 Count/Second | No | The maximum number of PutMedia requests per second that you can send for each feed, input, and stream combination in the current AWS Region. |
| Request rate for PutMedia, in a burst | Each supported Region: 5 Count | No | The maximum number of PutMedia requests that you can send in one burst for each feed, input, and stream combination in the current AWS Region. |
| Request rate for UpdateFeed | Each supported Region: 5 Count/Second | No | The maximum number of UpdateFeed requests per second that you can send in this account in the current AWS Region. |
| Request rate for UpdateFeed, in a burst | Each supported Region: 30 Count | No | The maximum number of UpdateFeed requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Total dictionaries per account | Each supported Region: 10 Count | Yes | The maximum number of dictionaries that you can create in this account in the current AWS Region. |
| Total feeds per account | Each supported Region: 1,000 Count | No | The maximum number of feeds that you can create in this account in the current AWS Region. |
| Total outputs per feed | Each supported Region: 30 Count | No | The maximum number of outputs that you can create per feed in this account in the current AWS Region. |

# Amazon MemoryDB endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | memory-db.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | memory-db-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | memory-db.us-east-1.amazonaws.com | HTTPS |
| | | memory-db-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | memory-db.us-west-1.amazonaws.com | HTTPS |
| | | memory-db-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | memory-db.us-west-2.amazonaws.com | HTTPS |
| | | memory-db-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | memory-db.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | memory-db.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | memory-db.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | memory-db.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | memory-db.ap-southeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | memory-db.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | memory-db.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | memory-db.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | memory-db.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | memory-db.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | memory-db.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | memory-db.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | memory-db.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | memory-db.eu-north-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | memory-db.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | memory-db.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | memory-db.us-gov-west-1.amazonaws.com | HTTPS | |

### Service quotas

| Resource | Default |
|---|---|
| Nodes per Region | 300 |
| Nodes per cluster per instance type | 90 |
| Nodes per shard | 6 |
| Parameter groups per Region | 150 |
| Subnet groups per Region | 150 |
| Subnets per subnet group | 20 |

## AWS Entity Resolution endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | entityresolution.us-east-2.amazonaws.com | HTTPS | |
| | | entityresolution.us-east-2.api.aws | HTTPS | |
| | | entityresolution-fips.us-east-2.amazonaws.com | HTTPS | |
| | | entityresolution-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | entityresolution.us-east-1.amazonaws.com | HTTPS | |
| | | entityresolution-fips.us-east-1.amazonaws.com | HTTPS | |
| | | entityresolution.us-east-1.api.aws | HTTPS | |
| | | entityresolution-fips.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | entityresolution.us-west-2.amazonaws.com | HTTPS | |
| | | entityresolution.us-west-2.api.aws | HTTPS | |
| | | entityresolution-fips.us-west-2.amazonaws.com | HTTPS | |
| | | entityresolution-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | entityresolution.af-south-1.amazonaws.com | HTTPS | |
| | | entityresolution.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | entityresolution.ap-northeast-2.amazonaws.com | HTTPS | |
| | | | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | entityresolution.ap-northeast-2.api.aws | |
| Asia Pacific (Singapore) | ap-southeast-1 | entityresolution.ap-southeast-1.amazonaws.com<br>entityresolution.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | entityresolution.ap-southeast-2.amazonaws.com<br>entityresolution.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | entityresolution.ap-northeast-1.amazonaws.com<br>entityresolution.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | entityresolution.ca-central-1.amazonaws.com<br>entityresolution.ca-central-1.api.aws<br>entityresolution-fips.ca-central-1.amazonaws.com<br>entityresolution-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | entityresolution.eu-central-1.amazonaws.com<br>entityresolution.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | entityresolution.eu-west-1.amazonaws.com<br>entityresolution.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | entityresolution.eu-west-2.amazonaws.com<br>entityresolution.eu-west-2.api.aws | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent ID mapping jobs | Each supported Region: 1 | No | The maximum number of ID mapping workflows that can be processed concurrently in the current AWS Region. |
| Concurrent matching jobs | Each supported Region: 1 | No | The maximum number of matching workflows that can be processed concurrently in the current AWS Region. |
| ID mapping workflows | Each supported Region: 10 | Yes | The maximum number of ID mapping workflows that you can create in this account in the current AWS Region. |
| ID namespaces | Each supported Region: 10 | Yes | The maximum number of ID namespaces that you can create in this account in the current AWS Region. |
| Matching workflows | Each supported Region: 10 | Yes | The maximum number of matching workflows that you can create in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GenerateMatchId API requests | Each supported Region: 10 | Yes | The maximum number of GenerateMatchId API requests per second |
| Rate of GetMatchId API requests | Each supported Region: 50 | Yes | The maximum number of GetMatchId API requests per second. |
| Records per machine learning-based matching workflow | Each supported Region: 150,000,000 | Yes | The maximum number of records that can be processed by a machine learning-based matching workflow in this account in AWS Regions of af-south-1, ap-northeast-2, eu-west-2. |
| Records per machine learning-based matching workflow | Each supported Region: 600,000,000 | Yes | The maximum number of records that can be processed by a machine learning-based matching workflow in this account in AWS Regions of ap-northeast-1, ap-southeast-1, ap-southeast-2, ca-central-1, eu-central-1, eu-west-1, us-east-1, us-east-2, us-west-2. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Records per provider ID mapping workflow | Each supported Region: 150,000,000 | Yes | The maximum number of records that can be processed for provider ID mapping in this account in AWS Regions of af-south-1, ap-northeast-2, eu-west-2. |
| Records per provider ID mapping workflow | Each supported Region: 250,000,000 | Yes | The maximum number of records that can be processed for provider ID mapping in this account in AWS Regions of ap-northeast-1, ap-southeast-1, ap-southeast-2, ca-central-1, eu-central-1, eu-west-1, us-east-1, us-east-2, us-west-2. |
| Records per provider service-based matching workflow | Each supported Region: 100,000,000 | Yes | The maximum number of records that can be processed by a provider service-based matching workflow in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Records per rule-based ID mapping workflow | Each supported Region: 1,000,000,000 | Yes | The maximum number of records that can be processed for rule-based ID mapping in this account in AWS Regions of ap-northeast-1, ap-southeast-1, ap-southeast-2, ca-central-1, eu-central-1, eu-west-1, us-east-1, us-east-2, us-west-2. |
| Records per rule-based ID mapping workflow | Each supported Region: 150,000,000 | Yes | The maximum number of records that can be processed for rule-based ID mapping in this account in the AWS Regions of af-south-1, ap-northeast-2, eu-west-2. |
| Records per rule-based matching workflow | Each supported Region: 100,000,000 | Yes | The maximum number of records that can be processed by a rule-based matching workflow in this account in the current AWS Region. |
| Schema mappings | Each supported Region: 50 | Yes | The maximum number of schema mappings that you can create in this account in the current AWS Region. |

For more information, see AWS Entity Resolution quotas in the *AWS Entity Resolution User Guide*.

# AWS End User Messaging endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

AWS End User Messaging includes the AWS End User Messaging SMS and Voice V2 API and AWS End User Messaging Social API.

## Service endpoints

### AWS End User Messaging SMS and Voice V2 API

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | sms-voice.us-east-2.amazonaws.com | HTTPS | |
| | | sms-voice-fips.us-east-2.api.aws | HTTPS | |
| | | sms-voice-fips.us-east-2.amazonaws.com | HTTPS | |
| | | sms-voice.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | sms-voice.us-east-1.amazonaws.com | HTTPS | |
| | | sms-voice-fips.us-east-1.api.aws | HTTPS | |
| | | sms-voice-fips.us-east-1.amazonaws.com | HTTPS | |
| | | sms-voice.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | sms-voice.us-west-1.amazonaws.com | HTTPS |
| | | sms-voice-fips.us-west-1.api.aws | HTTPS |
| | | sms-voice-fips.us-west-1.amazonaws.com | HTTPS |
| | | sms-voice.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | sms-voice.us-west-2.amazonaws.com | HTTPS |
| | | sms-voice-fips.us-west-2.api.aws | HTTPS |
| | | sms-voice-fips.us-west-2.amazonaws.com | HTTPS |
| | | sms-voice.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | sms-voice.af-south-1.amazonaws.com | HTTPS |
| | | sms-voice.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | sms-voice.ap-south-2.amazonaws.com | HTTPS |
| | | sms-voice.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | sms-voice.ap-southeast-3.amazonaws.com | HTTPS |
| | | sms-voice.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | sms-voice.ap-southeast-4.amazonaws.com | HTTPS |
| | | sms-voice.ap-southeast-4.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | sms-voice.ap-south-1.amazonaws.com | HTTPS |
| | | sms-voice.ap-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | sms-voice.ap-southeast-6.amazonaws.com | HTTPS |
| | | sms-voice.ap-southeast-6.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | sms-voice.ap-northeast-3.amazonaws.com | HTTPS |
| | | sms-voice.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | sms-voice.ap-northeast-2.amazonaws.com | HTTPS |
| | | sms-voice.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | sms-voice.ap-southeast-1.amazonaws.com | HTTPS |
| | | sms-voice.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | sms-voice.ap-southeast-2.amazonaws.com | HTTPS |
| | | sms-voice.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | sms-voice.ap-east-2.amazonaws.com | HTTPS |
| | | sms-voice.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | sms-voice.ap-northeast-1.amazonaws.com | HTTPS |
| | | sms-voice.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | sms-voice.ca-central-1.amazonaws.com | HTTPS | |
| | | sms-voice-fips.ca-central-1.api.aws | HTTPS | |
| | | sms-voice-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | sms-voice.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | sms-voice.ca-west-1.amazonaws.com | HTTPS | |
| | | sms-voice-fips.ca-west-1.api.aws | HTTPS | |
| | | sms-voice-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | sms-voice.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | sms-voice.eu-central-1.amazonaws.com | HTTPS | |
| | | sms-voice.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | sms-voice.eu-west-1.amazonaws.com | HTTPS | |
| | | sms-voice.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | sms-voice.eu-west-2.amazonaws.com | HTTPS | |
| | | sms-voice.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | sms-voice.eu-south-1.amazonaws.com | HTTPS | |
| | | sms-voice.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | sms-voice.eu-west-3.amazonaws.com | HTTPS | |
| | | sms-voice.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | sms-voice.eu-south-2.amazonaws.com | HTTPS | |
| | | sms-voice.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | sms-voice.eu-north-1.amazonaws.com | HTTPS |
| | | sms-voice.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | sms-voice.eu-central-2.amazonaws.com | HTTPS |
| | | sms-voice.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | sms-voice.il-central-1.amazonaws.com | HTTPS |
| | | sms-voice.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | sms-voice.mx-central-1.amazonaws.com | HTTPS |
| | | sms-voice.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | sms-voice.me-south-1.amazonaws.com | HTTPS |
| | | sms-voice.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | sms-voice.me-central-1.amazonaws.com | HTTPS |
| | | sms-voice.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | sms-voice.sa-east-1.amazonaws.com | HTTPS |
| | | sms-voice.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | sms-voice.us-gov-east-1.amazonaws.com | HTTPS |
| | | sms-voice-fips.us-gov-east-1.api.aws | HTTPS |
| | | sms-voice-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | sms-voice.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | sms-voice.us-gov-west-1.amazonaws.com | HTTPS | |
| | | sms-voice-fips.us-gov-west-1.api.aws | HTTPS | |
| | | sms-voice-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | sms-voice.us-gov-west-1.api.aws | HTTPS | |

## AWS End User Messaging Social API

### Region availability

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | social-messaging.us-east-1.amazonaws.com | Version 20 and later |
| | | social-messaging.us-east-1.api.aws | |
| | | social-messaging-fips.us-east-1.amazonaws.com | |
| | | social-messaging-fips.us-east-1.api.aws | |
| US East (Ohio) | us-east-2 | social-messaging.us-east-2.amazonaws.com | Version 20 and later |
| | | social-messaging.us-east-2.api.aws | |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| | | social-messaging-fips.us-east-2.amazonaws.com<br><br>social-messaging-fips.us-east-2.api.aws | |
| US West (N. California) | us-west-1 | social-messaging.us-west-1.amazonaws.com<br><br>social-messaging.us-west-1.api.aws<br><br>social-messaging-fips.us-west-1.amazonaws.com<br><br>social-messaging-fips.us-west-1.api.aws | Version 20 and later |
| US West (Oregon) | us-west-2 | social-messaging.us-west-2.amazonaws.com<br><br>social-messaging.us-west-2.api.aws<br><br>social-messaging-fips.us-west-2.amazonaws.com<br><br>social-messaging-fips.us-west-2.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Africa (Cape Town) | af-south-1 | social-messaging.af-south-1.amazonaws.com<br><br>social-messaging.af-south-1.api.aws | Version 20 and later |
| Asia Pacific (Hong Kong) | ap-east-1 | social-messaging.ap-east-1.amazonaws.com<br><br>social-messaging.ap-east-1.api.aws | Version 20 and later |
| Asia Pacific (Taipei) | ap-east-2 | social-messaging.ap-east-2.amazonaws.com<br><br>social-messaging.ap-east-2.api.aws | Version 20 and later |
| Asia Pacific (Mumbai) | ap-south-1 | social-messaging.ap-south-1.amazonaws.com<br><br>social-messaging.ap-south-1.api.aws | Version 20 and later |
| Asia Pacific (Hyderabad) | ap-south-2 | social-messaging.ap-south-2.amazonaws.com<br><br>social-messaging.ap-south-2.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | social-messaging.ap-northeast-1.amazonaws.com<br><br>social-messaging.ap-northeast-1.api.aws | Version 20 and later |
| Asia Pacific (Seoul) | ap-northeast-2 | social-messaging.ap-northeast-2.amazonaws.com<br><br>social-messaging.ap-northeast-2.api.aws | Version 20 and later |
| Asia Pacific (Osaka) | ap-northeast-3 | social-messaging.ap-northeast-3.amazonaws.com<br><br>social-messaging.ap-northeast-3.api.aws | Version 20 and later |
| Asia Pacific (Singapore) | ap-southeast-1 | social-messaging.ap-southeast-1.amazonaws.com<br><br>social-messaging.ap-southeast-1.api.aws | Version 20 and later |
| Asia Pacific (Sydney) | ap-southeast-2 | social-messaging.ap-southeast-2.amazonaws.com<br><br>social-messaging.ap-southeast-2.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | social-messaging.ap-southeast-3.amazonaws.com<br><br>social-messaging.ap-southeast-3.api.aws | Version 20 and later |
| Asia Pacific (Melbourne) | ap-southeast-4 | social-messaging.ap-southeast-4.amazonaws.com<br><br>social-messaging.ap-southeast-4.api.aws | Version 20 and later |
| Asia Pacific (Malaysia) | ap-southeast-5 | social-messaging.ap-southeast-5.amazonaws.com<br><br>social-messaging.ap-southeast-5.api.aws | Version 20 and later |
| Asia Pacific (New Zealand) | ap-southeast-6 | social-messaging.ap-southeast-6.amazonaws.com<br><br>social-messaging.ap-southeast-6.api.aws | Version 20 and later |
| Asia Pacific (Thailand) | ap-southeast-7 | social-messaging.ap-southeast-7.amazonaws.com<br><br>social-messaging.ap-southeast-7.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Canada (Central) | ca-central-1 | social-messaging.ca-central-1.amazonaws.com<br><br>social-messaging.ca-central-1.api.aws<br><br>social-messaging-fips.ca-central-1.amazonaws.com<br><br>social-messaging-fips.ca-central-1.api.aws | Version 20 and later |
| Canada West (Calgary) | ca-west-1 | social-messaging.ca-west-1.amazonaws.com<br><br>social-messaging.ca-west-1.api.aws<br><br>social-messaging-fips.ca-west-1.amazonaws.com<br><br>social-messaging-fips.ca-west-1.api.aws | Version 20 and later |
| Europe (Frankfurt) | eu-central-1 | social-messaging.eu-central-1.amazonaws.com<br><br>social-messaging.eu-central-1.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | social-messaging.eu-central-2.amazonaws.com<br><br>social-messaging.eu-central-2.api.aws | Version 20 and later |
| Europe (Stockholm) | eu-north-1 | social-messaging.eu-north-1.amazonaws.com<br><br>social-messaging.eu-north-1.api.aws | Version 20 and later |
| Europe (Milan) | eu-south-1 | social-messaging.eu-south-1.amazonaws.com<br><br>social-messaging.eu-south-1.api.aws | Version 20 and later |
| Europe (Spain) | eu-south-2 | social-messaging.eu-south-2.amazonaws.com<br><br>social-messaging.eu-south-2.api.aws | Version 20 and later |
| Europe (Ireland) | eu-west-1 | social-messaging.eu-west-1.amazonaws.com<br><br>social-messaging.eu-west-1.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Europe (London) | eu-west-2 | social-messaging.eu-west-2.amazonaws.com<br><br>social-messaging.eu-west-2.api.aws | Version 20 and later |
| Europe (Paris) | eu-west-3 | social-messaging.eu-west-3.amazonaws.com<br><br>social-messaging.eu-west-3.api.aws | Version 20 and later |
| Israel (Tel Aviv) | il-central-1 | social-messaging.il-central-1.amazonaws.com<br><br>social-messaging.il-central-1.api.aws | Version 20 and later |
| Middle East (Bahrain) | me-south-1 | social-messaging.me-south-1.amazonaws.com<br><br>social-messaging.me-south-1.api.aws | Version 20 and later |
| Middle East (UAE) | me-central-1 | social-messaging.me-central-1.amazonaws.com<br><br>social-messaging.me-central-1.api.aws | Version 20 and later |

| Region name | Region | Endpoint | WhatsApp API version |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | social-messaging.mx-central-1.amazonaws.com<br><br>social-messaging.mx-central-1.api.aws | Version 20 and later |
| South America (São Paulo) | sa-east-1 | social-messaging.sa-east-1.amazonaws.com<br><br>social-messaging.sa-east-1.api.aws | Version 20 and later |

# Service quotas

## AWS End User Messaging SMS and Voice Service Quotas

The following quotas apply to the SMS and MMS channel.

| Resource | Default quota | Eligible for increase |
|---|---|---|
| Spending threshold | USD $1.00 per account | Yes, but spending limits vary by region. You must specify the region(s) in which you require an increase. |
| Number of SMS messages that can be sent each second (*sending rate*) | Varies depending on destination country and originating phone number. | Yes, however, you might need to obtain a phone number that supports higher throughput. If you're unsure of which number type to use, contact Support or your AWS Account Manager for more information |

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| | | If you use an alphanumeric Sender ID to send messages, you might be able to increase your throughput rate. |
| Number of SMS and MMS messages that can be sent to a single recipient each second | 1 message per second | No |
| Number of Amazon SNS topics for two-way SMS | 100,000 per account | Yes |
| Number of Keywords for two-way SMS | 30 Keywords per number | Yes |
| Number of SMS, MMS, and Voice numbers | 25 per account | Yes |
| Number of dedicated phone numbers | 25 per account | Yes |
| Number of opt-out lists<br><br>Note: The required Default opt-out list counts against this quota. | 25 per account | Yes |
| Number of configuration sets | 25 per account | Yes |
| Number of event destinations | 5 per configuration set | No |
| Number of verified destinati on phone numbers while in SMS sandbox | 10 per account | No |
| Number of phone number pools | 25 per account | Yes |

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Number of origination identities that can be associated with a phone number pool | 100 per phone number pool | Yes |

The following quotas apply to SMS messages sent using 10DLC phone numbers. 10DLC numbers can only be used to send messages to recipients in the United States.

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Max 10DLC companies per AWS account | 25 | Yes |
| Max 10DLC campaigns per 10DLC brand | 10 | Yes |
| Max 10DLC numbers per 10DLC campaign | 49 | No |

The following quotas apply to protect configurations.

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Number of protect configura tions | 25 | No |

The following quotas apply to the voice channel.

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Number of voice messages that can be sent during a 24-hour period | If your account is in the sandbox: 20 messages | No |

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Number of voice messages that can be sent to a single recipient during a 24-hour period | 5 messages | No |
| Number of voice messages that can be sent per minute | If your account is in the sandbox: 5 calls per minute<br><br>If your account is out of the sandbox: 20 calls per minute | No |
| Number of voice messages that can be sent from a single originating phone number per second | 1 message per second | No |
| Voice message length | If your account is in the sandbox: 30 seconds<br><br>If your account is out of the sandbox: 5 minutes | No |

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Ability to send voice messages to international phone numbers | If your account is in the sandbox, you can send messages to recipients in only the following countries:<br>• Australia<br>• Canada<br>• Germany<br>• Hong Kong<br>• Israel<br>• Japan<br>• Mexico<br>• Singapore<br>• Sweden<br>• United States<br>• United Kingdom<br><br>If your account is out of the sandbox, you can send messages to recipients in any country.<br><br>**Note**<br>International calls are subject to additiona l fees, which vary by destination country or region. | No |

| Resource | Default quota | Eligible for increase |
| --- | --- | --- |
| Number of characters in a voice message | 3,000 billable characters, in words that are spoken<br><br>6,000 characters total, including billable characters and SSML tags | No |
| Number of configuration sets | 10,000 voice configuration sets | No |

For more information, see Quotas for AWS End User Messaging SMS in the *AWS End User Messaging SMS User Guide*.

### AWS End User Messaging Social Service Quotas

The following quotas apply to the Social messaging channel.

| Resource | Default |
| --- | --- |
| WhatsApp Business Account (WABA) | 25 per region |

For more information, see Quotas for AWS End User Messaging Social in the *AWS End User Messaging Social User Guide*.

## Amazon EMR endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | elasticmapreduce.us-east-2.amazonaws.com<br>elasticmapreduce-fips.us-east-2.amazonaws.com<br>elasticmapreduce.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | elasticmapreduce.us-east-1.amazonaws.com<br>elasticmapreduce-fips.us-east-1.amazonaws.com<br>elasticmapreduce.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | elasticmapreduce.us-west-1.amazonaws.com<br>elasticmapreduce-fips.us-west-1.amazonaws.com<br>elasticmapreduce.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | elasticmapreduce.us-west-2.amazonaws.com<br>elasticmapreduce-fips.us-west-2.amazonaws.com<br>elasticmapreduce.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | elasticmapreduce.af-south-1.amazonaws.com<br>elasticmapreduce.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific | ap-east-1 | elasticmapreduce.ap-east-1.amazonaws.com<br>elasticmapreduce.ap-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | elasticmapreduce.ap-south-2.amazonaws.com<br>elasticmapreduce.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | elasticmapreduce.ap-southeast-3.amazonaws.com<br>elasticmapreduce.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | elasticmapreduce.ap-southeast-5.amazonaws.com<br>elasticmapreduce.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | elasticmapreduce.ap-southeast-4.amazonaws.com<br>elasticmapreduce.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | elasticmapreduce.ap-south-1.amazonaws.com<br>elasticmapreduce.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | elasticmapreduce.ap-southeast-6.amazonaws.com<br>elasticmapreduce.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | elasticmapreduce.ap-northeast-3.amazonaws.com<br>elasticmapreduce.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | elasticmapreduce.ap-northeast-2.amazonaws.com<br>elasticmapreduce.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | elasticmapreduce.ap-southeast-1.amazonaws.com<br>elasticmapreduce.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | elasticmapreduce.ap-southeast-2.amazonaws.com<br>elasticmapreduce.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | elasticmapreduce.ap-east-2.amazonaws.com<br>elasticmapreduce.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | elasticmapreduce.ap-southeast-7.amazonaws.com<br>elasticmapreduce.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | elasticmapreduce.ap-northeast-1.amazonaws.com<br>elasticmapreduce.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | elasticmapreduce.ca-central-1.amazonaws.com<br>elasticmapreduce-fips.ca-central-1.amazonaws.com<br>elasticmapreduce.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | elasticmapreduce.ca-west-1.amazonaws.com | HTTPS |
| | | elasticmapreduce-fips.ca-west-1.amazonaws.com | HTTPS |
| | | elasticmapreduce.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | elasticmapreduce.eu-central-1.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | elasticmapreduce.eu-west-1.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | elasticmapreduce.eu-west-2.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | elasticmapreduce.eu-south-1.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | elasticmapreduce.eu-west-3.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | elasticmapreduce.eu-south-2.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | elasticmapreduce.eu-north-1.amazonaws.com | HTTPS |
| | | elasticmapreduce.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | elasticmapreduce.eu-central-2.amazonaws.com<br>elasticmapreduce.eu-central-2.api.aws | HTTPS<br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | elasticmapreduce.il-central-1.amazonaws.com<br>elasticmapreduce.il-central-1.api.aws | HTTPS<br>HTTPS | |
| Mexico (Central) | mx-central-1 | elasticmapreduce.mx-central-1.amazonaws.com<br>elasticmapreduce.mx-central-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | elasticmapreduce.me-south-1.amazonaws.com<br>elasticmapreduce.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | elasticmapreduce.me-central-1.amazonaws.com<br>elasticmapreduce.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | elasticmapreduce.sa-east-1.amazonaws.com<br>elasticmapreduce.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | elasticmapreduce.us-gov-east-1.amazonaws.com<br>elasticmapreduce.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | elasticmapreduce.us-gov-west-1.amazonaws.com<br><br>elasticmapreduce.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

If you specify the general endpoint (elasticmapreduce.amazonaws.com), Amazon EMR directs your request to an endpoint in the default Region. For accounts created on or after March 8, 2013, the default Region is us-west-2; for older accounts, the default Region is us-east-1.

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| EMR WAL workspaces | Each supported Region: 20 | No | The maximum number of Amazon EMR write-ahead log (EMR WAL) workspaces that can be used at the same time in this account in the current Region. |
| Replenishment rate of AddInstanceFleet calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the AddInstanceFleet bucket |
| Replenishment rate of AddInstanceGroups calls | Each supported Region: 0.2 | Yes | The Rate at which tokens are added to the AddInstanceGroups bucket |
| Replenishment rate of AddJobFlowSteps calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the AddJobFlowSteps bucket |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replenishment rate of AddTags calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the AddTags bucket |
| Replenishment rate of CancelJobRun calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the CancelJobRun bucket |
| Replenishment rate of CancelSteps calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the CancelSteps bucket |
| Replenishment rate of CreateManagedCluster calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the CreateManagedCluster bucket |
| Replenishment rate of CreateSecurityConfiguration calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the CreateSecurityConfiguration bucket |
| Replenishment rate of DeleteSecurityConfiguration calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the DeleteSecurityConfiguration bucket |
| Replenishment rate of DescribeCluster calls | Each supported Region: 1 | Yes | The Rate at which tokens are added to the DescribeCluster bucket |
| Replenishment rate of DescribeJobFlows calls | Each supported Region: 0.2 | Yes | The Rate at which tokens are added to the DescribeJobFlows bucket |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replenishment rate of DescribeJobRun calls | Each supported Region: 2 | Yes | The Rate at which tokens are added to the DescribeJobRun bucket |
| Replenishment rate of DescribeManagedCluster calls | Each supported Region: 2 | Yes | The Rate at which tokens are added to the DescribeManagedCluster bucket |
| Replenishment rate of DescribeSecurityConfiguration calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the DescribeSecurityConfiguration bucket |
| Replenishment rate of DescribeStep calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the DescribeStep bucket |
| Replenishment rate of GetAutoTerminationPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the GetAutoTermination Policy bucket |
| Replenishment rate of GetBlockPublicAccessConfiguration calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the GetBlockPublicAccessConfiguration bucket |
| Replenishment rate of GetClusterSessionCredentials calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the GetClusterSessionCredentials bucket |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replenishment rate of GetJobRunQueryResults calls | Each supported Region: 2 | Yes | The Rate at which tokens are added to the GetJobRunQueryResults bucket |
| Replenishment rate of GetManagedScalingPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the GetManagedScalingPolicy bucket |
| Replenishment rate of GetResourceDashboard calls | Each supported Region: 2 | Yes | The Rate at which tokens are added to the GetResourceDashboard bucket |
| Replenishment rate of ListBootstrapActions calls | Each supported Region: 1 | Yes | The Rate at which tokens are added to the ListBootstrapActions bucket |
| Replenishment rate of ListClusters calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListClusters bucket |
| Replenishment rate of ListInstanceFleets calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListInstanceFleets bucket |
| Replenishment rate of ListInstanceGroups calls | Each supported Region: 1 | Yes | The Rate at which tokens are added to the ListInstanceGroups bucket |
| Replenishment rate of ListInstances calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListInstances bucket |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replenishment rate of ListJobRuns calls | Each supported Region: 1 | Yes | The Rate at which tokens are added to the ListJobRuns bucket |
| Replenishment rate of ListManagedClusters calls | Each supported Region: 1 | Yes | The Rate at which tokens are added to the ListManagedClusters bucket |
| Replenishment rate of ListReleaseLabels calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListReleaseLabels bucket |
| Replenishment rate of ListSecurityConfigurations calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListSecurityConfigurations bucket |
| Replenishment rate of ListSteps calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListSteps bucket |
| Replenishment rate of ListSupportedInstanceTypes calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ListSupportedInstanceTypes bucket |
| Replenishment rate of ModifyCluster calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ModifyCluster bucket |
| Replenishment rate of ModifyInstanceFleet calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ModifyInstanceFleet bucket |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replenishment rate of ModifyInstanceGroups calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the ModifyInstanceGroups bucket |
| Replenishment rate of PutAutoScalingPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the PutAutoScalingPolicy bucket |
| Replenishment rate of PutAutoTerminationPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the PutAutoTermination Policy bucket |
| Replenishment rate of PutBlockPublicAccessConfiguration calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the PutBlockPublicAccessConfiguration bucket |
| Replenishment rate of PutManagedScalingPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the PutManagedScalingPolicy bucket |
| Replenishment rate of RemoveAutoScalingPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the RemoveAutoScalingPolicy bucket |
| Replenishment rate of RemoveAutoTerminationPolicy calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the RemoveAutoTerminationPolicy bucket |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replenishment rate of RemoveTags calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the RemoveTags bucket |
| Replenishment rate of RunJobFlow calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the RunJobFlow bucket |
| Replenishment rate of SetKeepJobFlowAliveWhenNoSteps calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the RemoveManagedScalingPolicy bucket. |
| Replenishment rate of SetKeepJobFlowAliveWhenNoSteps calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the SetKeepJobFlowAliveWhenNoSteps bucket |
| Replenishment rate of SetTerminationProtection calls | Each supported Region: 0.2 | Yes | The Rate at which tokens are added to the SetTerminationProtection bucket |
| Replenishment rate of SetVisibleToAllUsers calls | Each supported Region: 0.2 | Yes | The Rate at which tokens are added to the SetVisibleToAllUsers bucket |
| Replenishment rate of StartJobRun calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the StartJobRun bucket |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replenishment rate of TerminateJobFlows calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the TerminateJobFlows bucket |
| Replenishment rate of TerminateManagedCluster calls | Each supported Region: 0.5 | Yes | The Rate at which tokens are added to the TerminateManagedCluster bucket |
| The maximum number of API requests that you can make per second. | Each supported Region: 200 per second | Yes | The maximum number of requests per second that you can perform in this account in the current Region for all EMR operations. |
| The maximum number of AddInstanceFleet API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of AddInstanceFleet requests per second that you can perform in this account in the current Region. AddInstanceFleet adds an instance fleet to a running cluster. |
| The maximum number of AddInstanceGroups API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of AddInstanceGroups requests per second that you can perform in this account in the current Region. Adds one or more instance groups to a running cluster. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of AddJobFlowSteps API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of AddJobFlowSteps requests per second that you can perform in this account in the current Region. AddJobFlowSteps adds new steps to a running cluster. |
| The maximum number of AddTags API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of AddTags requests per second that you can perform in this account in the current Region. Adds tags to an Amazon EMR resource. |
| The maximum number of CancelJobRun API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of CancelJobRun requests per second that you can perform in this account in the current Region. CancelJobRun cancels a running job run on a managed cluster. |
| The maximum number of CancelSteps API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of CancelSteps requests per second that you can perform in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of CreateManagedCluster API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of CreateManagedCluster requests per second that you can perform in this account in the current Region. CreateManagedCluster creates and starts running a new managed cluster. |
| The maximum number of CreateSecurityConfiguration API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of CreateSecurityConfiguration requests per second that you can perform in this account in the current Region. CreateSecurityConfiguration creates a security configuration, which is stored in the service and can be specified when a cluster is created. |
| The maximum number of DeleteSecurityConfiguration API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of DeleteSecurityConfiguration requests per second that you can perform in this account in the current Region. DeleteSecurityConfiguration deletes a security configuration. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of DescribeCluster API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of DescribeCluster requests per second that you can perform in this account in the current Region. DescribeCluster provides cluster-level details including status, hardware and software configuration, VPC settings, and so on. |
| The maximum number of DescribeJobFlows API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of DescribeJobFlows requests per second that you can perform in this account in the current Region. This API is deprecated and will eventually be removed. We recommend you use ListClusters, DescribeCluster, ListSteps, ListInstanceGroups and ListBootstrapActions instead. |
| The maximum number of DescribeJobRun API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of DescribeJobRun requests per second that you can perform in this account in the current Region. DescribeJobRun provides job-run-level details including status. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of DescribeManagedCluster API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of DescribeManagedCluster requests per second that you can perform in this account in the current Region. DescribeManagedCluster provides managed-cluster-level details including status and configuration. |
| The maximum number of DescribeSecurityConfiguration API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of DescribeSecurityConfiguration requests per second that you can perform in this account in the current Region. DescribeSecurityConfiguration provides the details of a security configuration by returning the configuration JSON. |
| The maximum number of DescribeStep API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of DescribeStep requests per second that you can perform in this account in the current Region. DescribeStep provides more detail about the cluster step. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of GetAutoTerminationPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of GetAutoTerminationPolicy requests per second that you can perform in this account in the current Region. GetAutoTerminationPolicy returns the auto-termination policy for an Amazon EMR cluster. |
| The maximum number of GetBlockPublicAccessConfiguration API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of GetBlockPublicAccessConfiguration requests per second that you can perform in this account in the current Region. GetBlockPublicAccessConfiguration returns the Amazon EMR block public access configuration for the account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of GetClusterSessionCredentials API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of GetClusterSessionCredentials requests per second that you can perform in this account in the current Region. GetClusterSessionCredentials provides temporary, HTTP basic credentials that are associated with a given runtime IAM role and used by a cluster with fine-grained access control activated. |
| The maximum number of GetJobRunQueryResults API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of GetJobRunQueryResults requests per second that you can perform in this account in the current Region. GetJobRunQueryResults retrieves query results for a job run. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of GetManagedScalingPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of GetManagedScalingPolicy requests per second that you can perform in this account in the current Region. GetManagedScalingPolicy fetches the attached managed scaling policy for an Amazon EMR cluster. |
| The maximum number of GetResourceDashboard API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of GetResourceDashboard requests per second that you can perform in this account in the current Region. GetResourceDashboard retrieves resource dashboard information for a managed cluster. |
| The maximum number of ListBootstrapActions API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of ListBootstrapActions requests per second that you can perform in this account in the current Region. ListBootstrapActions provides information about the bootstrap actions associated with a cluster. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of ListClusters API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of ListClusters requests per second that you can perform in this account in the current Region. ListClusters provides the status of all clusters visible to this AWS account. |
| The maximum number of ListInstanceFleets API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of ListInstanceFleets requests per second that you can perform in this account in the current Region. ListInstanceFleets lists all available details about the instance fleets in a cluster. |
| The maximum number of ListInstanceGroups API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of ListInstanceGroups requests per second that you can perform in this account in the current Region. ListInstanceGroups provides all available details about the instance groups in a cluster. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of ListInstances API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of ListInstances requests per second that you can perform in this account in the current Region. ListInstances provides information for all active EC2 instances and EC2 instances terminated in the last 30 days, up to a maximum of 2,000. |
| The maximum number of ListJobRuns API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of ListJobRuns requests per second that you can perform in this account in the current Region. ListJobRuns provides the status of all job runs on a managed cluster. |
| The maximum number of ListManagedClusters API requests that you can make per second. | Each supported Region: 20 per second | Yes | The maximum number of ListManagedClusters requests per second that you can perform in this account in the current Region. ListManagedClusters provides the status of all managed clusters visible to this AWS account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of ListReleaseLabels API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of ListReleaseLabels requests per second that you can perform in this account in the current Region. ListReleaseLabels retrieves release labels of Amazon EMR services in the region. |
| The maximum number of ListSecurityConfigurations API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of ListSecurityConfigurations requests per second that you can perform in this account in the current Region. ListSecurityConfigurations lists all the security configurations visible to this account, providing their creation dates and times, and their names. |
| The maximum number of ListSteps API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of ListSteps requests per second that you can perform in this account in the current Region. ListSteps provides a list of steps for the cluster in reverse order. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of ListSupportedInstanceTypes API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of ListSupportedInstanceTypes requests per second that you can perform in this account in the current Region. ListSupportedInstanceTypes lists the instance types that Amazon EMR supports. |
| The maximum number of ModifyCluster API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of ModifyCluster requests per second that you can perform in this account in the current Region. ModifyCluster modifies the number of steps that can be executed concurrently for the cluster specified using ClusterID. |
| The maximum number of ModifyInstanceFleet API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of ModifyInstanceFleet requests per second that you can perform in this account in the current Region. ModifyInstanceFleet modifies the target On-Demand and target Spot capacities for the instance fleet. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of ModifyInstanceGroups API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of ModifyInstanceGroups requests per second that you can perform in this account in the current Region. ModifyInstanceGroups modifies the number of nodes and configuration settings of an instance group. |
| The maximum number of PutAutoScalingPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of PutAutoScalingPolicy requests per second that you can perform in this account in the current Region. PutAutoScalingPolicy creates or updates an automatic scaling policy for a core instance group or task instance group in an Amazon EMR cluster. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of PutAutoTerminationPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of PutAutoTerminationPolicy requests per second that you can perform in this account in the current Region. PutAutoTerminationPolicy creates or updates an auto-termination policy for an Amazon EMR cluster. |
| The maximum number of PutBlockPublicAccessConfiguration API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of PutBlockPublicAccessConfiguration requests per second that you can perform in this account in the current Region. PutBlockPublicAccessConfiguration creates or updates an Amazon EMR block public access configuration. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of PutManagedScalingPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of PutManagedScalingPolicy requests per second that you can perform in this account in the current Region. PutManagedScalingPolicy creates or updates a managed scaling policy for an Amazon EMR cluster. |
| The maximum number of RemoveAutoScalingPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of RemoveAutoScalingPolicy requests per second that you can perform in this account in the current Region. RemoveAutoScalingPolicy removes an automatic scaling policy from a specified instance group within an EMR cluster. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of RemoveAutoTerminationPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of RemoveAutoTerminationPolicy requests per second that you can perform in this account in the current Region. RemoveAutoTerminationPolicy removes an auto-termination policy from an Amazon EMR cluster. |
| The maximum number of RemoveManagedScalingPolicy API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of RemoveManagedScalingPolicy requests per second that you can perform in this account in the current Region. PutManagedScalingPolicy removes a managed scaling policy from a specified Amazon EMR cluster. |
| The maximum number of RemoveTags API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of RemoveTags requests per second that you can perform in this account in the current Region. Removes tags from an Amazon EMR resource. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of RunJobFlow API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of RunJobFlow requests per second that you can perform in this account in the current Region. RunJobFlow creates and starts running a new cluster (job flow). |
| The maximum number of SetKeepJobFlowAliveWhenNoSteps API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of SetKeepJobFlowAliveWhenNoSteps per second that you can perform in this account in the current Region |
| The maximum number of SetTerminationProtection API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of SetTerminationProtection requests per second that you can perform in this account in the current Region. SetTerminationProtection locks a cluster (job flow) so the EC2 instances in the cluster cannot be terminated. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of SetVisibleToAllUsers API requests that you can make per second. | Each supported Region: 5 per second | Yes | The maximum number of SetVisibleToAllUsers requests per second that you can perform in this account in the current Region. Sets the VisibleToAllUsers value, which determines whether the cluster is visible to all IAM users of the AWS account associated with the cluster. |
| The maximum number of StartJobRun API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of StartJobRun requests per second that you can perform in this account in the current Region. StartJobRun submits a new job run to a managed cluster. |
| The maximum number of TerminateJobFlows API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of TerminateJobFlows requests per second that you can perform in this account in the current Region. TerminateJobFlows shuts a list of clusters (job flows) down. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| The maximum number of TerminateManagedCluster API requests that you can make per second. | Each supported Region: 10 per second | Yes | The maximum number of TerminateManagedCluster requests per second that you can perform in this account in the current Region. TerminateManagedCluster shuts down a managed cluster. |
| The maximum number of active clusters can be run at the same time | Each supported Region: 500 | Yes | The maximum number of active clusters can be run at the same time. |
| The maximum number of active instances per instance group | Each supported Region: 2,000 | Yes | The maximum number of active instances per instance group. |
| The maximum rate at which your bucket replenishes for all EMR operations. | Each supported Region: 50 | Yes | The maximum rate at which your bucket replenishes for all EMR operations. |
| Write-ahead logs (EMR WAL) | Each supported Region: 20,000 | No | The maximum number of Amazon EMR write-ahead logs (EMR WAL) that can be used at the same time in this account in the current Region. |

Amazon EMR throttles the following API requests for each AWS account on a per-Region basis. For more information about how throttling is applied, see API Request Throttling in the *Amazon EC2 API Reference*.

# Amazon EventBridge endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | events.us-east-2.amazonaws.com | HTTPS |
| | | events-fips.us-east-2.api.aws | HTTPS |
| | | events-fips.us-east-2.amazonaws.com | HTTPS |
| | | events.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | events.us-east-1.amazonaws.com | HTTPS |
| | | events-fips.us-east-1.api.aws | HTTPS |
| | | events-fips.us-east-1.amazonaws.com | HTTPS |
| | | events.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | events.us-west-1.amazonaws.com | HTTPS |
| | | events-fips.us-west-1.api.aws | HTTPS |
| | | events-fips.us-west-1.amazonaws.com | HTTPS |
| | | events.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | events.us-west-2.amazonaws.com | HTTPS | |
| | | events-fips.us-west-2.api.aws | HTTPS | |
| | | events-fips.us-west-2.amazonaws.com | HTTPS | |
| | | events.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | events.af-south-1.amazonaws.com | HTTPS | |
| | | events.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | events.ap-east-1.amazonaws.com | HTTPS | |
| | | events.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | events.ap-south-2.amazonaws.com | HTTPS | |
| | | events.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | events.ap-southeast-3.amazonaws.com | HTTPS | |
| | | events.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | events.ap-southeast-5.amazonaws.com | HTTPS | |
| | | events.ap-southeast-5.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | events.ap-southeast-4.amazonaws.com | HTTPS | |
| | | events.ap-southeast-4.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | events.ap-south-1.amazonaws.com<br>events.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | events.ap-southeast-6.amazonaws.com<br>events.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | events.ap-northeast-3.amazonaws.com<br>events.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | events.ap-northeast-2.amazonaws.com<br>events.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | events.ap-southeast-1.amazonaws.com<br>events.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | events.ap-southeast-2.amazonaws.com<br>events.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | events.ap-east-2.amazonaws.com<br>events.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | events.ap-southeast-7.amazonaws.com<br>events.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | events.ap-northeast-1.amazonaws.com | HTTPS |
| | | events.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | events.ca-central-1.amazonaws.com | HTTPS |
| | | events-fips.ca-central-1.api.aws | HTTPS |
| | | events-fips.ca-central-1.amazonaws.com | HTTPS |
| | | events.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | events.ca-west-1.amazonaws.com | HTTPS |
| | | events-fips.ca-west-1.api.aws | HTTPS |
| | | events-fips.ca-west-1.amazonaws.com | HTTPS |
| | | events.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | events.eu-central-1.amazonaws.com | HTTPS |
| | | events.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | events.eu-west-1.amazonaws.com | HTTPS |
| | | events.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | events.eu-west-2.amazonaws.com | HTTPS |
| | | events.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | events.eu-south-1.amazonaws.com | HTTPS |
| | | events.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | events.eu-west-3.amazonaws.com | HTTPS |
| | | events.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | events.eu-south-2.amazonaws.com | HTTPS |
| | | events.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | events.eu-north-1.amazonaws.com | HTTPS |
| | | events.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | events.eu-central-2.amazonaws.com | HTTPS |
| | | events.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | events.il-central-1.amazonaws.com | HTTPS |
| | | events.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | events.mx-central-1.amazonaws.com | HTTPS |
| | | events.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | events.me-south-1.amazonaws.com | HTTPS |
| | | events.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | events.me-central-1.amazonaws.com | HTTPS |
| | | events.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | events.sa-east-1.amazonaws.com | HTTPS |
| | | events.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | events.us-gov-east-1.amazonaws.com | HTTPS |
| | | events.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | events.us-gov-west-1.amazonaws.com<br><br>events.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Api destinations | Each supported Region: 3,000 | Yes | The maximum number of API destinations per account per Region. |
| Connections | Each supported Region: 3,000 | Yes | The maximum number of connections per account per Region. |
| CreateEndpoint throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of requests per second for CreateEndpoint API. Additional requests are throttled. |
| DeleteEndpoint throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of requests per second for DeleteEndpoint API. Additional requests are throttled. |
| Endpoints | Each supported Region: 100 | Yes | The maximum number of endpoints per account per Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Event bus policy size | Each supported Region: 10,240 | Yes | Maximum policy size, in characters. This policy size increases each time you grant access to another account. You can see your current policy and its size by using the DescribeEventBus API. |
| Event buses | Each supported Region: 100 | Yes | Maximum event buses per account. |
| Event pattern size | Each supported Region: 2,048 | Yes | Maximum size of an event pattern, in characters. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Invocations throttle limit in transactions per second | us-east-1: 18,750 per second<br><br>us-east-2: 4,500 per second<br><br>us-west-1: 2,250 per second<br><br>us-west-2: 18,750 per second<br><br>ap-east-1: 1,100 per second<br><br>ap-northeast-1: 2,250 per second<br><br>ap-northeast-2: 1,100 per second<br><br>ap-south-1: 1,100 per second<br><br>ap-southeast-1: 2,250 per second<br><br>ap-southeast-2: 2,250 per second<br><br>ca-central-1: 1,100 per second<br><br>eu-central-1: 4,500 per second | Yes | An invocation is an event matching a rule and being sent on to the rules targets. After the limit is reached, the invocations are throttled; that is, they still happen but they are delayed. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-north-1: 1,100 per second<br><br>eu-south-2: 4,500 per second<br><br>eu-west-1: 18,750 per second<br><br>eu-west-2: 2,250 per second<br><br>eu-west-3: 1,100 per second<br><br>sa-east-1: 1,100 per second<br><br>Each of the other supported Regions: 750 per second | | |
| Number of rules | af-south-1: 100<br><br>eu-south-1: 100<br><br>Each of the other supported Regions: 300 | Yes | Maximum number of rules an account can have per event bus |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| PutEvents throttle limit in transactions per second | us-east-1: 10,000 per second<br><br>us-east-2: 2,400 per second<br><br>us-west-1: 1,200 per second<br><br>us-west-2: 10,000 per second<br><br>ap-east-1: 600 per second<br><br>ap-northeast-1: 1,200 per second<br><br>ap-northeast-2: 600 per second<br><br>ap-south-1: 600 per second<br><br>ap-southeast-1: 1,200 per second<br><br>ap-southeast-2: 1,200 per second<br><br>ca-central-1: 600 per second<br><br>eu-central-1: 2,400 per second | Yes | Maximum number of requests per second for PutEvents API. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-north-1: 600 per second<br><br>eu-south-2: 2,400 per second<br><br>eu-west-1: 10,000 per second<br><br>eu-west-2: 1,200 per second<br><br>eu-west-3: 600 per second<br><br>sa-east-1: 600 per second<br><br>Each of the other supported Regions: 400 per second | | |
| Rate of invocations per API destination | Each supported Region: 300 per second | Yes | The maximum number of invocations per second to send to each API destination endpoint per account per Region. Once the quota is met, future invocations to that API endpoint are throttled. The invocations will still occur, but are delayed. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Targets per rule | Each supported Region: 5 | No | Maximum number of targets that can be associated with a rule |
| Throttle limit in transactions per second for control plane APIs | Each supported Region: 50 per second | No | Maximum number of requests per second for EventBridge control plane API operations. Additional requests are throttled. |
| UpdateEndpoint throttle limit in transactions per second | Each supported Region: 5 per second | No | The maximum number of requests per second for UpdateEndpoint API. Additional requests are throttled. |

## Amazon EventBridge Pipes endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | pipes.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | pipes.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | pipes.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | pipes.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | pipes.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | pipes.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | pipes.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | pipes.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | pipes.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | pipes.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | pipes.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | pipes.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | pipes.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | pipes.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | pipes.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | pipes.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | pipes.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | pipes.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | pipes.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | pipes.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | pipes.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | pipes.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | pipes.eu-central-2.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | pipes.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | pipes.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | pipes.sa-east-1.amazonaws.com | HTTPS | |

### FIPS service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | pipes-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | pipes-fips.us-west-2.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | pipes-fips.us-west-1.amazonaws.com | HTTPS |
| US East (Ohio) | us-east-2 | pipes-fips.us-east-2.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | pipes-fips.ca-central-1.amazonaws.com | HTTPS |

## Service quotas

EventBridge Pipes has the following quotas. If you have requirements for higher maximum limits, contact support.

| Resource | Regions | Default limit |
|---|---|---|
| Concurrent pipe executions per account | • AWS GovCloud (US-West)<br>• AWS GovCloud (US-East)<br>• China (Ningxia)<br>• China (Beijing)<br>• Asia Pacific (Osaka) | 1000 |

| Resource | Regions | Default limit |
| --- | --- | --- |
| | • Africa (Cape Town)<br>• Europe (Milan)<br>• US East (Ohio)<br>• Europe (Frankfurt)<br>• US West (N. California)<br>• Europe (London)<br>• Asia Pacific (Sydney)<br>• Asia Pacific (Tokyo)<br>• Asia Pacific (Singapore)<br>• Canada (Central)<br>• Europe (Paris)<br>• Europe (Stockholm)<br>• South America (São Paulo)<br>• Asia Pacific (Seoul)<br>• Asia Pacific (Mumbai)<br>• Asia Pacific (Hong Kong)<br>• Middle East (Bahrain)<br>• China (Ningxia)<br>• China (Beijing)<br>• Asia Pacific (Osaka)<br>• Africa (Cape Town)<br>• Europe (Milan) | |
| Concurrent pipe executions per account | • US East (N. Virginia)<br>• US West (Oregon)<br>• Europe (Ireland) | 3000 |
| Pipes per account | All | 1000 |

# Amazon EventBridge Scheduler endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | scheduler.us-east-2.amazonaws.com | HTTPS |
| | | scheduler-fips.us-east-2.amazonaws.com | HTTPS |
| | | scheduler-fips.us-east-2.api.aws | HTTPS |
| | | scheduler.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | scheduler.us-east-1.amazonaws.com | HTTPS |
| | | scheduler-fips.us-east-1.amazonaws.com | HTTPS |
| | | scheduler-fips.us-east-1.api.aws | HTTPS |
| | | scheduler.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | scheduler.us-west-1.amazonaws.com | HTTPS |
| | | scheduler-fips.us-west-1.amazonaws.com | HTTPS |
| | | scheduler-fips.us-west-1.api.aws | HTTPS |
| | | scheduler.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | scheduler.us-west-2.amazonaws.com | HTTPS |
| | | scheduler-fips.us-west-2.amazonaws.com | HTTPS |
| | | scheduler-fips.us-west-2.api.aws | HTTPS |
| | | scheduler.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | scheduler.af-south-1.amazonaws.com | HTTPS |
| | | scheduler.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | scheduler.ap-east-1.amazonaws.com | HTTPS |
| | | scheduler.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | scheduler.ap-south-2.amazonaws.com | HTTPS |
| | | scheduler.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | scheduler.ap-southeast-3.amazonaws.com | HTTPS |
| | | scheduler.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | scheduler.ap-southeast-5.amazonaws.com | HTTPS |
| | | scheduler.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | scheduler.ap-southeast-4.amazonaws.com | HTTPS |
| | | scheduler.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | scheduler.ap-south-1.amazonaws.com<br>scheduler.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | scheduler.ap-southeast-6.amazonaws.com<br>scheduler.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | scheduler.ap-northeast-3.amazonaws.com<br>scheduler.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | scheduler.ap-northeast-2.amazonaws.com<br>scheduler.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | scheduler.ap-southeast-1.amazonaws.com<br>scheduler.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | scheduler.ap-southeast-2.amazonaws.com<br>scheduler.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | scheduler.ap-east-2.amazonaws.com<br>scheduler.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | scheduler.ap-southeast-7.amazonaws.com<br>scheduler.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | scheduler.ap-northeast-1.amazonaws.com | HTTPS |
| | | scheduler.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | scheduler.ca-central-1.amazonaws.com | HTTPS |
| | | scheduler-fips.ca-central-1.amazonaws.com | HTTPS |
| | | scheduler-fips.ca-central-1.api.aws | HTTPS |
| | | scheduler.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | scheduler.ca-west-1.amazonaws.com | HTTPS |
| | | scheduler-fips.ca-west-1.amazonaws.com | HTTPS |
| | | scheduler-fips.ca-west-1.api.aws | HTTPS |
| | | scheduler.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | scheduler.eu-central-1.amazonaws.com | HTTPS |
| | | scheduler.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | scheduler.eu-west-1.amazonaws.com | HTTPS |
| | | scheduler.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | scheduler.eu-west-2.amazonaws.com | HTTPS |
| | | scheduler.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | scheduler.eu-south-1.amazonaws.com | HTTPS |
| | | scheduler.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | scheduler.eu-west-3.amazonaws.com | HTTPS |
| | | scheduler.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | scheduler.eu-south-2.amazonaws.com | HTTPS | |
| | | scheduler.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | scheduler.eu-north-1.amazonaws.com | HTTPS | |
| | | scheduler.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | scheduler.eu-central-2.amazonaws.com | HTTPS | |
| | | scheduler.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | scheduler.il-central-1.amazonaws.com | HTTPS | |
| | | scheduler.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | scheduler.mx-central-1.amazonaws.com | HTTPS | |
| | | scheduler.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | scheduler.me-south-1.amazonaws.com | HTTPS | |
| | | scheduler.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | scheduler.me-central-1.amazonaws.com | HTTPS | |
| | | scheduler.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | scheduler.sa-east-1.amazonaws.com | HTTPS | |
| | | scheduler.sa-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | scheduler.us-gov-east-1.amazonaws.com | HTTPS | |
| | | scheduler-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | scheduler-fips.us-gov-east-1.api.aws | HTTPS | |
| | | scheduler.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | scheduler.us-gov-west-1.amazonaws.com | HTTPS | |
| | | scheduler-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | scheduler-fips.us-gov-west-1.api.aws | HTTPS | |
| | | scheduler.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CreateSchedule request rate | us-east-1: 5,000<br>us-east-2: 5,000<br>us-west-2: 5,000<br>ap-northeast-1: 5,000<br>ap-south-1: 5,000<br>ap-southeast-1: 5,000<br>ap-southeast-2: 5,000 | Yes | Maximum CreateSchedule requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. This is adjustable to tens of thousands of requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-central-1: 5,000<br><br>eu-west-1: 5,000<br><br>eu-west-2: 5,000<br><br>sa-east-1: 5,000<br><br>Each of the other supported Regions: 250 | | |
| CreateScheduleGroup request rate | Each supported Region: 10 | Yes | Maximum CreateScheduleGroup requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteSchedule request rate | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>eu-west-1: 1,000<br>eu-west-2: 1,000<br>sa-east-1: 1,000<br>Each of the other supported Regions: 250 | Yes | Maximum DeleteSchedule requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. This is adjustable to tens of thousands of requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteScheduleGroup request rate | Each supported Region: 10 | Yes | Maximum DeleteScheduleGroup requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetSchedule request rate | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>eu-west-1: 1,000<br>eu-west-2: 1,000<br>sa-east-1: 1,000<br>Each of the other supported Regions: 250 | Yes | Maximum GetSchedule requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. This is adjustable to tens of thousands of requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetScheduleGroup request rate | Each supported Region: 10 | Yes | Maximum GetScheduleGroup requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Invocations throttle limit in transactions per second | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>eu-west-1: 1,000<br>eu-west-2: 1,000<br>sa-east-1: 1,000<br>Each of the other supported Regions: 500 | Yes | An invocation is a schedule payload being delivered to the defined target. After the limit is reached, the invocations are throttled; that is, they still happen but they are delayed. This is adjustable to tens of thousands of transactions per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListScheduleGroups request rate | Each supported Region: 10 | Yes | Maximum ListScheduleGroups requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. |
| ListSchedules request rate | Each supported Region: 50 | Yes | Maximum ListSchedules requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. |
| ListTagsForResource request rate | Each supported Region: 10 | Yes | Lists the tags associated with the Scheduler resource. |
| Number of schedule groups | Each supported Region: 500 | Yes | Maximum number of schedule groups per region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of schedules | Each supported Region: 10,000,000 | Yes | The maximum number of schedules per region. This quota includes one-time schedules that have completed running. We recommend configuring your schedules to automatically delete after completion using the ActionAfterCompletion feature. This is adjustable to billions of schedules. |
| TagResource request rate | Each supported Region: 1 | Yes | Assigns one or more tags (key-value pairs) to the specified Scheduler resource. |
| UntagResource request rate | Each supported Region: 1 | Yes | Removes one or more tags from the specified Scheduler resource. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateSchedule request rate | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-south-1: 1,000<br>ap-southeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-central-1: 1,000<br>eu-west-1: 1,000<br>eu-west-2: 1,000<br>sa-east-1: 1,000<br>Each of the other supported Regions: 250 | Yes | Maximum UpdateSchedule requests per second. When you reach this quota, EventBridge Scheduler rejects requests for this operation for the remainder of the interval. This is adjustable to tens of thousands of requests per second. |

# Amazon EventBridge Schemas endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | schemas.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | schemas.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | schemas.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | schemas.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | schemas.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | schemas.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | schemas.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | schemas.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | schemas.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | schemas.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | schemas.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | schemas.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | schemas.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | schemas.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | schemas.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | schemas.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | schemas.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | schemas.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | schemas.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | schemas.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | schemas.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | schemas.eu-central-2.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | schemas.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | schemas.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | schemas.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | schemas.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | schemas.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DiscoveredSchemas | Each supported Region: 200 | Yes | The maximum number of schemas for a discovered schema registry that you can create in the current region |
| Discoverers | Each supported Region: 10 | Yes | The maximum number of discoverers that you can create in the current region. |
| Registries | Each supported Region: 10 | Yes | The maximum number of registries that you can create in the current region. |
| SchemaVersions | Each supported Region: 100 | Yes | The maximum number of versions per schema that you can create in the current region. |
| Schemas | Each supported Region: 100 | Yes | The maximum number of schemas per registry that you can create in the |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | current region. (Except Discovered Schema Registry) |

# Amazon Elastic VMware Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | evs.us-east-2.amazonaws.com | HTTPS | |
| | | evs-fips.us-east-2.api.aws | HTTPS | |
| | | evs-fips.us-east-2.amazonaws.com | HTTPS | |
| | | evs.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | evs.us-east-1.amazonaws.com | HTTPS | |
| | | evs-fips.us-east-1.api.aws | HTTPS | |
| | | evs-fips.us-east-1.amazonaws.com | HTTPS | |
| | | evs.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | evs.us-west-1.amazonaws.com | HTTPS |
| | | evs-fips.us-west-1.api.aws | HTTPS |
| | | evs-fips.us-west-1.amazonaws.com | HTTPS |
| | | evs.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | evs.us-west-2.amazonaws.com | HTTPS |
| | | evs-fips.us-west-2.api.aws | HTTPS |
| | | evs-fips.us-west-2.amazonaws.com | HTTPS |
| | | evs.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | evs.ap-south-2.amazonaws.com | HTTPS |
| | | evs.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | evs.ap-southeast-5.amazonaws.com | HTTPS |
| | | evs.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | evs.ap-south-1.amazonaws.com | HTTPS |
| | | evs.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | evs.ap-northeast-3.amazonaws.com | HTTPS |
| | | evs.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | evs.ap-northeast-2.amazonaws.com | HTTPS |
| | | evs.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | evs.ap-southeast-1.amazonaws.com<br>evs.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | evs.ap-southeast-2.amazonaws.com<br>evs.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | evs.ap-east-2.amazonaws.com<br>evs.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | evs.ap-northeast-1.amazonaws.com<br>evs.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | evs.ca-central-1.amazonaws.com<br>evs.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | evs.ca-west-1.amazonaws.com<br>evs.ca-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | evs.eu-central-1.amazonaws.com<br>evs.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | evs.eu-west-1.amazonaws.com<br>evs.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | evs.eu-west-2.amazonaws.com<br>evs.eu-west-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | evs.eu-south-1.amazonaws.com | HTTPS | |
| | | evs.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | evs.eu-west-3.amazonaws.com | HTTPS | |
| | | evs.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | evs.eu-south-2.amazonaws.com | HTTPS | |
| | | evs.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | evs.eu-north-1.amazonaws.com | HTTPS | |
| | | evs.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | evs.eu-central-2.amazonaws.com | HTTPS | |
| | | evs.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | evs.il-central-1.amazonaws.com | HTTPS | |
| | | evs.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | evs.mx-central-1.amazonaws.com | HTTPS | |
| | | evs.mx-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | evs.sa-east-1.amazonaws.com | HTTPS | |
| | | evs.sa-east-1.api.aws | HTTPS | |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Environment count per AWS account | Each supported Region: 3 | Yes | The maximum number of EVS environments that can be created in this account in the current Region. |
| Host count per EVS environment | Each supported Region: 5 | Yes | Maximum number of Hosts that can be provisioned within a single EVS environment |

These service quotas are listed under **Amazon Elastic VMware Service** in the Service Quotas console. To request a quota increase for values that are shown as adjustable, see Requesting a quota increase in the *Service Quotas User Guide*.

**Important**

Ensure that your EC2 Running On-Demand Standard Instance quota reflects the number of vCPUs that you need for all of the EC2 instances that you will use on Amazon EVS. Each i4i.metal instance uses 128 vCPUs. For information about increasing EC2 service quotas, see Request an increase in the *Amazon EC2 User Guide*.

**Note**

If you plan to use EC2 Dedicated Hosts for your Amazon EVS environment, ensure that your EC2 Dedicated i4i Hosts quota reflects the number of Dedicated Hosts that you intend to use for a desired Region. For information about increasing EC2 service quotas, see Request an increase in the *Amazon EC2 User Guide*.

**Note**

Amazon CloudWatch collects AWS usage metrics for Amazon EVS resources that have quotas (environment and hosts). For more information, see CloudWatch Usage Metrics in the *Amazon CloudWatch User Guide*.

**Note**

The default values are the initial quotas set by AWS. These default values are separate from the actual applied quota values and maximum possible service quotas. For more information, see Terminology in Service Quotas in the *Service Quotas User Guide*.

# Amazon FinSpace quotas

The following are the service quotas for this service. Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Attribute sets | Each supported Region: 100 | Yes | The maximum number of attribute sets that can exist in a FinSpace environment. |
| Clusters per user | Each supported Region: 1 | No | The maximum number of FinSpace Spark clusters that can be active for each user. |
| Concurrent changesets processing | Each supported Region: 10 | Yes | The maximum number of concurrent changeset |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | s that can be processing per FinSpace environment |
| Concurrent data views processing | Each supported Region: 10 | Yes | The maximum number of concurrently running data views processing per FinSpace environment. |
| Concurrent dataview version processing | Each supported Region: 10 | Yes | The maximum number of concurrent Managed kdb dataview version processing. |
| Controlled vocabularies and categories | Each supported Region: 100 | Yes | The maximum combined number of controlled vocabularies and categories per FinSpace environment. |
| Data views per dataset | Each supported Region: 3 | Yes | The maximum number of data views that can be created per dataset. |
| Datasets | Each supported Region: 1,500 | Yes | The maximum number of datasets that can exist in a FinSpace environment. |
| Datasets per permission group | Each supported Region: 1,500 | Yes | The maximum number of datasets assigned per permission group. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Environments | Each supported Region: 2 | Yes | The maximum number of FinSpace environments you can create per AWS account. |
| File size per changeset | Each supported Region: 50 Gigabytes | No | The maximum file size of any single file in a changeset. |
| Files per changeset | Each supported Region: 100,000 | No | The maximum number of files in a single changeset. |
| Managed kdb cluster users | Each supported Region: 1,000 | Yes | The maximum number of cluster users per environment. |
| Managed kdb clusters | Each supported Region: 10 | Yes | The maximum number of clusters allowed per environment. |
| Managed kdb concurrent changeset ingestions | Each supported Region: 10 | Yes | The maximum number of concurrent changeset ingestions allowed per environment. |
| Managed kdb database cluster cache size | Each supported Region: 7,730 Gigabytes | Yes | The maximum amount of database cluster cache allowed per environment. |
| Managed kdb databases | Each supported Region: 1,500 | Yes | The maximum number of databases allowed per environment. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Managed kdb dataviews | Each supported Region: 4,500 | Yes | The maximum number of Managed kdb dataviews per environment. |
| Managed kdb multi-AZ clusters | Each supported Region: 1 | Yes | The maximum number of multi-AZ clusters per environment. |
| Managed kdb savedown storage | Each supported Region: 17,179 Gigabytes | Yes | The maximum amount of savedown storage allowed per environment. |
| Managed kdb scaling groups | Each supported Region: 10 | Yes | The maximum number of Managed kdb scaling groups per environment. |
| Managed kdb single-AZ clusters | Each supported Region: 5 | Yes | The maximum number of single-AZ clusters per environment. |
| Managed kdb volume read mounts | Each supported Region: 5 | Yes | The maximum number of read mounts per Managed kdb volume per environment. |
| Managed kdb volume storage | Each supported Region: 7,730 Gigabytes | Yes | The maximum amount of storage for Managed kdb volumes per environment. |
| Managed kdb volume write mounts | Each supported Region: 5 | Yes | The maximum number of write mounts per Managed kdb volume per environment. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Managed kdb volumes | Each supported Region: 5 | Yes | The maximum number of Managed kdb volumes per environment. |
| Notebook storage | Each supported Region: 10 Gigabytes | No | The maximum amount of EFS storage per user notebook environment. |
| Permission groups | Each supported Region: 20 | Yes | The maximum number of permission groups per FinSpace environment. |
| Total kdb environments | Each supported Region: 1 | No | The maximum number of environments per AWS account. |
| Users | Each supported Region: 5 | Yes | The maximum number of users that can exist in a FinSpace environment. |
| kx.s.16xlarge nodes | Each supported Region: 0 | Yes | The maximum number of kx.s.16xlarge nodes allowed per environment. |
| kx.s.2xlarge nodes | Each supported Region: 5 | Yes | The maximum number of kx.s.2xlarge nodes allowed per environment. |
| kx.s.32xlarge nodes | Each supported Region: 0 | Yes | The maximum number of kx.s.32xlarge nodes allowed per environment. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| kx.s.4xlarge nodes | Each supported Region: 1 | Yes | The maximum number of kx.s.4xlarge nodes allowed per environment. |
| kx.s.8xlarge nodes | Each supported Region: 1 | Yes | The maximum number of kx.s.8xlarge nodes allowed per environment. |
| kx.s.large nodes | Each supported Region: 5 | Yes | The maximum number of kx.s.large nodes allowed per environment. |
| kx.s.xlarge nodes | Each supported Region: 5 | Yes | The maximum number of kx.s.xlarge nodes allowed per environment. |
| kx.sg.16xlarge scaling group nodes | Each supported Region: 0 | Yes | The maximum number of kx.sg.16xlarge Managed kdb scaling group nodes per environment. |
| kx.sg.2xlarge scaling group nodes | Each supported Region: 1 | Yes | The maximum number of kx.sg.2xlarge Managed kdb scaling group nodes per environment. |
| kx.sg.32xlarge scaling group nodes | Each supported Region: 0 | Yes | The maximum number of kx.sg.32xlarge Managed kdb scaling group nodes per environment. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| kx.sg.4xlarge scaling group nodes | Each supported Region: 1 | Yes | The maximum number of kx.sg.4xlarge Managed kdb scaling group nodes per environment. |
| kx.sg.8xlarge scaling group nodes | Each supported Region: 1 | Yes | The maximum number of kx.sg.8xlarge Managed kdb scaling group nodes per environment. |
| kx.sg.large scaling group nodes | Each supported Region: 1 | Yes | The maximum number of kx.sg.large Managed kdb scaling group nodes per environment. |
| kx.sg.xlarge scaling group nodes | Each supported Region: 1 | Yes | The maximum number of kx.sg.xlarge Managed kdb scaling group nodes per environment. |
| kx.sg1.16xlarge scaling group nodes | Each supported Region: 0 | Yes | The maximum number of kx.sg1.16xlarge Managed kdb scaling group nodes per environment. |
| kx.sg1.24xlarge scaling group nodes | Each supported Region: 0 | Yes | The maximum number of kx.sg1.24xlarge Managed kdb scaling group nodes per environment. |

# AWS Fault Injection Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | fis.us-east-2.amazonaws.com<br>fis-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | fis.us-east-1.amazonaws.com<br>fis-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | fis.us-west-1.amazonaws.com<br>fis-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | fis.us-west-2.amazonaws.com<br>fis-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | fis.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | fis.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | fis.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | fis.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | fis.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | fis.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | fis.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | fis.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | fis.ca-central-1.amazonaws.com<br>fis-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | fis.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | fis.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | fis.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | fis.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | fis.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | fis.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | fis.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | fis.eu-central-2.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | fis.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | fis.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | fis.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | fis.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | fis.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Action duration in hours | Each supported Region: 12 | No | The maximum number of hours allowed to run one action in this account in the current Region. |
| Action minutes per multi-Region strongly consistent (MRSC) global table for aws:dynamodb:global-table-pause-replication action | Each supported Region: 5,040 | No | The maximum cumulative action minutes per MRSC global table that aws:dynamodb:global-table-pause-replication can target in a 7-day rolling window. |
| Actions per experiment template | Each supported Region: 20 | No | The maximum number of actions that you can create in an experiment template in this account in the current Region. |
| Active experiments | Each supported Region: 5 | No | The maximum number of active experiments that you can run simultane |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | ously in this account in the current Region. |
| Completed experiment data retention in days | Each supported Region: 120 | No | The maximum number of days allowed for AWS FIS to retain data about completed experiments in this account in the current Region. |
| Experiment duration in hours | Each supported Region: 12 | No | The maximum number of hours allowed to run one experiment in this account in the current Region. |
| Experiment templates | Each supported Region: 500 | No | The maximum number of experiment templates that you can create in this account in the current Region. |
| Maximum number of Managed Prefix Lists in aws:network:route-table-disrupt-cross-region-connectivity | Each supported Region: 15 | No | The maximum number of Managed Prefix Lists that aws:network:route-table-disrupt-cross-region-connectivity will allow, per action. |
| Maximum number of Route Tables in aws:network:route-table-disrupt-cross-region-connectivity | Each supported Region: 10 | No | The maximum number of Route Tables that aws:network:route-table-disrupt-cross-region-connectivity will allow, per action. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of routes in aws:network:route-table-disrupt-cross-region-connectivity | Each supported Region: 200 | No | The maximum number of routes that aws:network:route-table-disrupt-cross-region-connectivity will allow, per action. |
| Parallel actions per experiment | Each supported Region: 10 | No | The maximum number of actions that you can run in parallel in an experiment in this account in the current Region. |
| Stop conditions per experiment template | Each supported Region: 5 | No | The maximum number of stop conditions that you can add to an experiment template in this account in the current Region. |
| Target Aurora DSQL clusters for aws:dsql:cluster-connection-failure action. | Each supported Region: 100 | No | The maximum number of Aurora DSQL clusters that aws:dsql:cluster-connection-failure can target when you identify targets using tags, per experiment. |
| Target Auto Scaling groups for aws:ec2:asg-insufficient-instance-capacity-error | Each supported Region: 500 | No | The maximum number of Auto Scaling groups that aws:ec2:asg-insufficient-instance-capacity-error can target when you identify targets using tags, per experiment. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Target Buckets for aws:s3:bucket-pause-replication | Each supported Region: 25 | No | The maximum number of S3 Buckets that aws:s3:bucket-pause-replication can target when you identify targets using tags, per experiment. |
| Target Clusters for aws:ecs:drain-container-instances | Each supported Region: 100 | No | The maximum number of Clusters that aws:ecs:drain-container-instances can target when you identify targets using tags, per experiment. |
| Target Clusters for aws:rds:failover-db-cluster | Each supported Region: 160 | No | The maximum number of Clusters that aws:rds:failover-db-cluster can target when you identify targets using tags, per experiment. |
| Target DBInstances for aws:rds:reboot-db-instances | Each supported Region: 130 | No | The maximum number of DBInstances that aws:rds:reboot-db-instances can target when you identify targets using tags, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target Instances for aws:ec2:reboot-instances | Each supported Region: 600 | No | The maximum number of Instances that aws:ec2:reboot-instances can target when you identify targets using tags, per experiment. |
| Target Instances for aws:ec2:stop-instances | Each supported Region: 400 | No | The maximum number of Instances that aws:ec2:stop-instances can target when you identify targets using tags, per experiment. |
| Target Instances for aws:ec2:terminate-instances | Each supported Region: 300 | No | The maximum number of Instances that aws:ec2:terminate-instances can target when you identify targets using tags, per experiment. |
| Target Instances for aws:ssm:send-command | Each supported Region: 200 | No | The maximum number of Instances that aws:ssm:send-command can target when you identify targets using tags, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target Kinesis data streams for aws:kinesis:stream-expired-iterator-exception action. | Each supported Region: 280 | No | The maximum number of Kinesis data streams that aws:kinesis:stream-expired-iterator-exception can target when you identify targets using tags, per experiment. |
| Target Kinesis data streams for aws:kinesis:stream-provisioned-throughput-exception action. | Each supported Region: 280 | No | The maximum number of Kinesis data streams that aws:kinesis:stream-provisioned-throughput-exception can target when you identify targets using tags, per experiment. |
| Target ManagedResources for aws:arc:start-zonal-autoshift action. | Each supported Region: 200 | No | The maximum number of managed resources that aws:arc:start-zonal-autoshift can target when you identify targets using tags, per experiment. |
| Target Nodegroups for aws:eks:terminate-nodegroup-instances | Each supported Region: 100 | No | The maximum number of Nodegroups that aws:eks:terminate-nodegroup-instances can target when you identify targets using tags, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target Pods for aws:eks:pod-cpu-stress | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-cpu-stress can target when you identify targets using parameters, per experiment. |
| Target Pods for aws:eks:pod-delete | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-delete can target when you identify targets using parameters, per experiment. |
| Target Pods for aws:eks:pod-io-stress | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-io-stress can target when you identify targets using parameters, per experiment. |
| Target Pods for aws:eks:pod-memory-stress | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-memory-stress can target when you identify targets using parameters, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target Pods for aws:eks:pod-network-blackhole-port | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-network-blackhole-port can target when you identify targets using parameters, per experiment. |
| Target Pods for aws:eks:pod-network-latency | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-network-latency can target when you identify targets using parameters, per experiment. |
| Target Pods for aws:eks:pod-network-packet-loss | Each supported Region: 1,000 | No | The maximum number of Pods that aws:eks:pod-network-packet-loss can target when you identify targets using parameters, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target ReplicationGroups for aws:elasticache:replicationgroup-interrupt-az-power | Each supported Region: 50 | No | The maximum number of ReplicationGroups that aws:elasticache:replicationgroup-interrupt-az-power can target per experiment. A daily limit applies to targeting ReplicationGroups. For more information, visit: https://docs.aws.amazon.com/fis/latest/userguide/fis-quotas.html. |
| Target SpotInstances for aws:ec2:send-spot-instance-interruptions | Each supported Region: 100 | No | The maximum number of SpotInstances that aws:ec2:send-spot-instance-interruptions can target when you identify targets using tags, per experiment. |
| Target Subnets for aws:network:disrupt-connectivity | Each supported Region: 100 | No | The maximum number of Subnets that aws:network:disrupt-connectivity can target when you identify targets using tags, per experiment. Quotas above 5 apply only to parameter scope:all. If you require a higher quota for another scope type, contact customer support. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Target Subnets for aws:network:route-table-disrupt-cross-region-connectivity | Each supported Region: 50 | No | The maximum number of Subnets that aws:network:route-table-disrupt-cross-region-connectivity can target when you identify targets using tags, per experiment. |
| Target Tasks for aws:ecs:stop-task | Each supported Region: 500 | No | The maximum number of Tasks that aws:ecs:stop-task can target when you identify targets using tags, per experiment. |
| Target Tasks for aws:ecs:task-cpu-stress | Each supported Region: 200 | No | The maximum number of Tasks that aws:ecs:task-cpu-stress can target when you identify targets using tags/parameters, per experiment. |
| Target Tasks for aws:ecs:task-io-stress | Each supported Region: 200 | No | The maximum number of Tasks that aws:ecs:task-io-stress can target when you identify targets using tags/parameters, per experiment. |
| Target Tasks for aws:ecs:task-kill-process | Each supported Region: 200 | No | The maximum number of Tasks that aws:ecs:task-kill-process can target when you identify targets using tags/parameters, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target Tasks for aws:ecs:task-network-blackhole-port | Each supported Region: 200 | No | The maximum number of Tasks that aws:ecs:task-network-blackhole-port can target when you identify targets using tags/parameters, per experiment. |
| Target Tasks for aws:ecs:task-network-latency | Each supported Region: 200 | No | The maximum number of Tasks that aws:ecs:task-network-latency can target when you identify targets using tags/parameters, per experiment. |
| Target Tasks for aws:ecs:task-network-packet-loss | Each supported Region: 200 | No | The maximum number of Tasks that aws:ecs:task-network-packet-loss can target when you identify targets using tags/parameters, per experiment. |
| Target TransitGateways for aws:network:transit-gateway-disrupt-cross-region-connectivity | Each supported Region: 50 | No | The maximum number of Transit Gateways that aws:network:transit-gateway-disrupt-cross-region-connectivity can target when you identify targets using tags, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target Virtual Interfaces for aws:directconnect:virtual-interface | Each supported Region: 5 | No | The maximum number of Virtual Interfaces that aws:directconnect:virtual-interface can target when you identify targets using tags, per experiment. |
| Target Volumes for aws:ebs:pause-volume-io | Each supported Region: 160 | No | The maximum number of Volumes that aws:ebs:pause-volume-io can target when you identify targets using tags, per experiment. |
| Target account configurations per experiment template | Each supported Region: 40 | No | The maximum number of target account configurations that you can create for an experiment template in this account in the current Region. |
| Target functions for aws:lambda:invocation-add-delay action. | Each supported Region: 140 | No | The maximum number of Lambda functions that aws:lambda:invocation-add-delay can target when you identify targets using tags, per experiment. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Target functions for aws:lambda:invocation-error action. | Each supported Region: 140 | No | The maximum number of Lambda functions that aws:lambda:invocation-error can target when you identify targets using tags, per experiment. |
| Target functions for aws:lambda:invocation-http-integration-response action. | Each supported Region: 140 | No | The maximum number of Lambda functions that aws:lambda:invocation-http-integration-response can target when you identify targets using tags, per experiment. |
| Target multi-Region clusters for aws:memorydb:multi-region-cluster-pause-replication action. | Each supported Region: 50 | No | The maximum number of MemoryDB multi-Region clusters that aws:memorydb:multi-region-cluster-pause-replication can target when you identify targets using tags, per experiment. If you require a higher quota, contact customer support. |
| Target tables for aws:dynamodb:global-table-pause-replication action | Each supported Region: 60 | No | The maximum number of global tables that aws:dynamodb:global-table-pause-replication can target, per experiment. |

# Amazon Data Firehose endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | firehose.us-east-2.amazonaws.com | HTTPS |
| | | firehose-fips.us-east-2.api.aws | HTTPS |
| | | firehose.us-east-2.api.aws | HTTPS |
| | | firehose-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | firehose.us-east-1.amazonaws.com | HTTPS |
| | | firehose-fips.us-east-1.amazonaws.com | HTTPS |
| | | firehose.us-east-1.api.aws | HTTPS |
| | | firehose-fips.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | firehose.us-west-1.amazonaws.com | HTTPS |
| | | firehose.us-west-1.api.aws | HTTPS |
| | | firehose-fips.us-west-1.amazonaws.com | HTTPS |
| | | firehose-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | firehose.us-west-2.amazonaws.com | HTTPS |
| | | firehose-fips.us-west-2.api.aws | HTTPS |
| | | firehose.us-west-2.api.aws | HTTPS |
| | | firehose-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | firehose.af-south-1.amazonaws.com | HTTPS |
| | | firehose.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | firehose.ap-east-1.amazonaws.com | HTTPS |
| | | firehose.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | firehose.ap-south-2.amazonaws.com | HTTPS |
| | | firehose.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | firehose.ap-southeast-3.amazonaws.com | HTTPS |
| | | firehose.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | firehose.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | firehose.ap-southeast-4.amazonaws.com | HTTPS |
| | | firehose.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | firehose.ap-south-1.amazonaws.com<br>firehose.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | firehose.ap-southeast-6.amazonaws.com<br>firehose.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | firehose.ap-northeast-3.amazonaws.com<br>firehose.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | firehose.ap-northeast-2.amazonaws.com<br>firehose.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | firehose.ap-southeast-1.amazonaws.com<br>firehose.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | firehose.ap-southeast-2.amazonaws.com<br>firehose.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | firehose.ap-east-2.amazonaws.com<br>firehose.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | firehose.ap-southeast-7.amazonaws.com<br>firehose.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | firehose.ap-northeast-1.amazonaws.com | HTTPS |
| | | firehose.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | firehose.ca-central-1.amazonaws.com | HTTPS |
| | | firehose.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | firehose.ca-west-1.amazonaws.com | HTTPS |
| | | firehose.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | firehose.eu-central-1.amazonaws.com | HTTPS |
| | | firehose.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | firehose.eu-west-1.amazonaws.com | HTTPS |
| | | firehose.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | firehose.eu-west-2.amazonaws.com | HTTPS |
| | | firehose.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | firehose.eu-south-1.amazonaws.com | HTTPS |
| | | firehose.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | firehose.eu-west-3.amazonaws.com | HTTPS |
| | | firehose.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | firehose.eu-south-2.amazonaws.com | HTTPS |
| | | firehose.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | firehose.eu-north-1.amazonaws.com | HTTPS |
| | | firehose.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | firehose.eu-central-2.amazonaws.com | HTTPS |
| | | firehose.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | firehose.il-central-1.amazonaws.com | HTTPS |
| | | firehose.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | firehose.mx-central-1.amazonaws.com | HTTPS |
| | | firehose.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | firehose.me-south-1.amazonaws.com | HTTPS |
| | | firehose.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | firehose.me-central-1.amazonaws.com | HTTPS |
| | | firehose.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | firehose.sa-east-1.amazonaws.com | HTTPS |
| | | firehose.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | firehose.us-gov-east-1.amazonaws.com | HTTPS |
| | | firehose-fips.us-gov-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | firehose.us-gov-west-1.amazonaws.com | HTTPS |
| | | firehose-fips.us-gov-west-1.amazonaws.com | HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Delivery streams | us-east-1: 5,000<br>us-east-2: 5,000<br>us-west-1: 2,000<br>us-west-2: 5,000<br>ap-northeast-1: 5,000<br>ap-northeast-2: 2,000<br>ap-northeast-3: 500<br>ap-south-1: 2,000<br>ap-southeast-1: 2,000<br>ap-southeast-2: 2,000<br>ca-central-1: 2,000<br>eu-central-1: 2,000<br>eu-north-1: 500<br>eu-south-1: 500<br>eu-west-1: 5,000 | Yes | The maximum number of delivery streams you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-west-2: 2,000<br><br>eu-west-3: 500<br><br>sa-east-1: 500<br><br>Each of the other supported Regions: 100 | | |
| Dynamic Partitions | Each supported Region: 500 | No | The maximum number of dynamic partitions for a delivery stream in the current Region. |
| Rate of CreateDeliveryStream requests | Each supported Region: 5 per second | No | The maximum number of CreateDeliveryStream requests you can make per second in this account in the current Region. |
| Rate of DeleteDeliveryStream requests | Each supported Region: 5 per second | No | The maximum number of DeleteDeliveryStream requests you can make per second in this account in the current Region. |
| Rate of DescribeDeliveryStream requests | Each supported Region: 5 | No | The maximum number of DescribeDeliveryStream requests you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListDeliveryStream requests | Each supported Region: 5 per second | No | The maximum number of ListDeliveryStream requests you can make per second in this account in the current Region. |
| Rate of ListTagsForDeliveryStream requests | Each supported Region: 5 per second | No | The maximum number of ListTagsForDeliveryStream requests you can make per second in this account in the current Region. |
| Rate of Put requests | us-east-1: 2,000<br>us-west-2: 2,000<br>eu-west-1: 2,000<br>Each of the other supported Regions: 1,000 | No | The maximum number of combined PutRecord and PutRecordBatch requests per second for a delivery stream in the current Region. |
| Rate of StartDeliveryStreamEncryption requests | Each supported Region: 5 per second | No | The maximum number of StartDeliveryStreamEncryption requests you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StopDeliveryStreamEncryption requests | Each supported Region: 5 per second | No | The maximum number of StopDeliveryStream Encryption requests you can make per second in this account in the current Region. |
| Rate of TagDeliveryStream requests | Each supported Region: 5 per second | No | The maximum number of TagDeliveryStream requests you can make per second in this account in the current Region. |
| Rate of UntagDeliveryStream requests | Each supported Region: 5 per second | No | The maximum number of UntagDeliveryStream requests you can make per second in this account in the current Region. |
| Rate of UpdateDestination requests | Each supported Region: 5 per second | No | The maximum number of UpdateDestination requests you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of data | us-east-1: 5<br><br>us-west-2: 5<br><br>eu-west-1: 5<br><br>Each of the other supported Regions: 1 | No | The maximum capacity in mebibyte per second for a delivery stream in the current Region. |
| Rate of records | us-east-1: 500,000<br><br>us-west-2: 500,000<br><br>eu-west-1: 500,000<br><br>Each of the other supported Regions: 100,000 | No | The maximum capacity in records per second for a delivery stream in the current Region. |

For more information, see Amazon Data Firehose Quotas in the *Amazon Data Firehose Developer Guide*.

## AWS Firewall Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | fms.us-east-2.amazonaws.com<br>fms-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | fms.us-east-1.amazonaws.com<br>fms-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | fms.us-west-1.amazonaws.com<br>fms-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | fms.us-west-2.amazonaws.com<br>fms-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | fms.af-south-1.amazonaws.com<br>fms-fips.af-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | fms.ap-east-1.amazonaws.com<br>fms-fips.ap-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | fms.ap-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | fms.ap-southeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | fms.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | fms.ap-southeast-4.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | fms.ap-south-1.amazonaws.com<br>fms-fips.ap-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | fms.ap-southeast-6.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | fms.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | fms.ap-northeast-2.amazonaws.com<br>fms-fips.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | fms.ap-southeast-1.amazonaws.com<br>fms-fips.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | fms.ap-southeast-2.amazonaws.com<br>fms-fips.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | fms.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | fms.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | fms.ap-northeast-1.amazonaws.com<br>fms-fips.ap-northeast-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | fms.ca-central-1.amazonaws.com<br>fms-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | fms.ca-west-1.amazonaws.com<br>fms-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | fms.eu-central-1.amazonaws.com<br>fms-fips.eu-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | fms.eu-west-1.amazonaws.com<br>fms-fips.eu-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | fms.eu-west-2.amazonaws.com<br>fms-fips.eu-west-2.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Milan) | eu-south-1 | fms.eu-south-1.amazonaws.com<br>fms-fips.eu-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | fms.eu-west-3.amazonaws.com<br>fms-fips.eu-west-3.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | fms.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | fms.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | fms.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | fms.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | fms.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | fms.me-south-1.amazonaws.com<br>fms-fips.me-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | fms.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | fms.sa-east-1.amazonaws.com<br>fms-fips.sa-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | fms.us-gov-east-1.amazonaws.com<br>fms-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | fms.us-gov-west-1.amazonaws.com<br>fms-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| AWS WAF Classic rule groups per AWS WAF Classic policy | Each supported Region: 2 | No | The maximum number of AWS WAF Classic rule groups that you can use in a Firewall Manager AWS WAF Classic policy. |
| Accounts per Firewall Manager admin | Each supported Region: 10,000 | Yes | The maximum number of accounts that you can be managed by Firewall Manager under a single FMS admin. |
| Admins per organization in Firewall Manager | Each supported Region: 10 | No | The maximum number of admin accounts that you can onboard under a single organization in Firewall Manager |
| Amazon VPC instances in scope of a common security group policy | Each supported Region: 100 | Yes | The maximum number of Amazon VPC instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | that you can have in scope per Firewall Manager common security group policy per account. This number represents the combined count of VPCs that you own and VPCs that are shared with you. |
| Applications per application list | Each supported Region: 50 | Yes | The maximum number of applications that you can define in an application list. |
| Audit security groups per security group content audit policy | Each supported Region: 1 | Yes | The maximum number of audit security groups that you can use in a Firewall Manager content audit security group policy. |
| Custom managed application lists for rules that allow all traffic | Each supported Region: 1 | Yes | The maximum number of managed application lists for rules that allow all traffic in a Firewall Manager content audit security group policy. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Custom managed application lists in any content audit security group policy setting | Each supported Region: 1 | Yes | The maximum number of custom managed application lists that you can use in any setting in a Firewall Manager content audit security group policy. |
| Custom managed application lists per account | Each supported Region: 10 | Yes | The maximum number of custom managed application lists that you can define for an account. |
| Custom managed protocol lists in any content audit security group policy setting | Each supported Region: 1 | Yes | The maximum number of custom managed protocol lists that you can use in any setting in a Firewall Manager content audit security group policy. |
| Custom managed protocol lists per account | Each supported Region: 10 | Yes | The maximum number of custom managed protocol lists that you can define for an account. |
| Explicitly included or excluded accounts per policy per Region | Each supported Region: 200 | Yes | The maximum number of accounts per Region that you can explicitly include in scope or explicitly exclude from scope for a Firewall Manager policy. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Firewall Manager policies per organization per Region | Each supported Region: 50 | Yes | The maximum number of Firewall Manager policies for any pair of Region and organization in AWS Organizations. |
| IPV4 CIDRs for a Network Firewall policy | Each supported Region: 50 | Yes | The maximum number of IPV4 CIDR ranges that you can provide in a single Firewall Manager Network Firewall policy, for use in firewall endpoint management. |
| Inbound/outbound rules per network ACL policy | Each supported Region: 5 | Yes | The maximum number of inbound/outbound rules that you can use in a Firewall Manager Network ACL policy. |
| Organizational units in scope per policy per Region | Each supported Region: 20 | Yes | The maximum number of organizational units that can be in scope of a Firewall Manager policy for any Region. |
| Partner rule groups per AWS WAF policy | Each supported Region: 1 | Yes | The maximum number of partner rule groups that can be applied to a single policy, limited by Web Application Firewall WebACL limits. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Primary security groups per common security group policy | Each supported Region: 3 | Yes | The maximum number of primary security groups that you can use in a Firewall Manager common security group policy. |
| Protocols per protocol list | Each supported Region: 5 | Yes | The maximum number of protocols that you can define in a protocol list. |
| Resource sets per Firewall Manager admin account | Each supported Region: 20 | Yes | The maximum number of resource sets that you can create under a single Firewall Manager admin account. |
| Resources per resource set | Each supported Region: 100 | Yes | The maximum number of resources that a single resource set can contain. |
| Route 53 Resolver DNS Firewall rule groups per DNS Firewall policy | Each supported Region: 2 | Yes | The maximum number of Route 53 Resolver DNS Firewall rule groups that you can use in a Firewall Manager DNS Firewall policy. |
| Rule groups per AWS WAF policy | Each supported Region: 50 | Yes | The maximum number of rule groups that you can use in a Firewall Manager AWS WAF policy. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Stateful rule group capacity per Network Firewall policy | Each supported Region: 30,000 | Yes | The maximum number of rules that can be included in a stateful rule group within a Firewall Manager Network Firewall policy. |
| Stateful rule groups per Network Firewall policy | Each supported Region: 20 | No | The maximum number of stateful rule groups that you can use in a Firewall Manager Network Firewall policy. |
| Stateless rule group capacity per Network Firewall policy | Each supported Region: 30,000 | No | The maximum number of rules that can be included in a stateless rule group within a Firewall Manager Network Firewall policy. |
| Stateless rule groups per Network Firewall policy | Each supported Region: 20 | No | The maximum number of stateless rule groups that you can use in a Firewall Manager Network Firewall policy. |
| Tags to include or exclude resources per policy | Each supported Region: 8 | Yes | The maximum number of tags that you can use to include or exclude resources for a Firewall Manager policy. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| VPCs that a single Network Firewall policy can automatically remediate | Each supported Region: 1,000 | No | The maximum number of VPCs that a single Firewall Manager Network Firewall policy can automatically remediate. |
| Web ACL capacity units (WCU) used in an AWS WAF policy | Each supported Region: 5,000 | No | The maximum combined number of web ACL capacity units (WCU) for all of the rule groups used in a Firewall Manager AWS WAF policy. The WCU usage for a rule group is fixed by the rule group owner at creation time. |

For more information, see AWS Firewall Manager quotas in the *AWS Firewall Manager Developer Guide*.

## Amazon Forecast endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Amazon Forecast

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | forecast.us-east-2.amazonaws.com<br>forecast-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | forecast.us-east-1.amazonaws.com<br>forecast-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | forecast.us-west-2.amazonaws.com<br>forecast-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | forecast.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | forecast.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | forecast.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | forecast.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | forecast.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | forecast.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | forecast.eu-west-1.amazonaws.com | HTTPS | |

## Amazon Forecast Query

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | forecastquery.us-east-2.amazonaws.com<br>forecastquery-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | forecastquery.us-east-1.amazonaws.com<br>forecastquery-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | forecastquery.us-west-2.amazonaws.com<br>forecastquery-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | forecastquery.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | forecastquery.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific | ap-southeast-1 | forecastquery.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Singapore) | | | | |
| Asia Pacific (Sydney) | ap-southeast-2 | forecastquery.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | forecastquery.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | forecastquery.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | forecastquery.eu-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum cumulative size of all files in your Amazon S3 bucket | Each supported Region: 30 Gigabytes | Yes | The maximum cumulative size of all your files in your Amazon S3 bucket in GB |
| Maximum forecast horizon | Each supported Region: 500 | No | The maximum prediction length (lesser of 500 data points or 1/3 of the target time series dataset length) |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of Explainabilities | Each supported Region: 1,000 | No | Maximum number of Explainabilities that you can have in your Amazon Forecast account |
| Maximum number of Explainability exports | Each supported Region: 1,000 | No | Maximum number of Explainability exports that you can have in your Amazon Forecast account |
| Maximum number of backtest windows | Each supported Region: 5 | No | The maximum Back-Test Window Count (RecipeParameters) |
| Maximum number of columns in a related time series dataset | Each supported Region: 25 | No | The maximum number of columns that can be in the RELATED_TIME_SERIES dataset |
| Maximum number of columns in a target time series dataset | Each supported Region: 13 | No | The maximum number of columns that can be in the TARGET_TIME_SERIES dataset |
| Maximum number of columns in an item metadata dataset | Each supported Region: 10 | No | The maximum number of columns that can be in the ITEM_METADATA dataset |
| Maximum number of dataset groups | Each supported Region: 500 | Yes | The maximum number of dataset groups that you can have in your Amazon Forecast account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of dataset import jobs | Each supported Region: 1,000 | Yes | The maximum number of dataset imports that you can have in your Amazon Forecast account |
| Maximum number of datasets | Each supported Region: 1,500 | Yes | The maximum number of datasets that you can have in your Amazon Forecast account |
| Maximum number of datasets in a dataset group | Each supported Region: 3 | No | The maximum number of datasets that can be in a dataset group |
| Maximum number of files in your Amazon S3 bucket | Each supported Region: 10,000 | No | The maximum number of files that you can have in your Amazon S3 bucket |
| Maximum number of forecast export jobs | Each supported Region: 1,000 | Yes | The maximum number of forecast exports that you can have in your Amazon Forecast account |
| Maximum number of forecasts | Each supported Region: 100 | Yes | The maximum number of forecasts that you can have in your Amazon Forecast account |
| Maximum number of predictor backtest export jobs | Each supported Region: 1,000 | Yes | The maximum number of predictor backtest exports that you can have in your Amazon Forecast account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of predictors | Each supported Region: 500 | Yes | The maximum number of predictors that you can have in your Amazon Forecast account |
| Maximum number of rows in a dataset | ap-south-1: 1,000,000,000<br><br>Each of the other supported Regions: 3,000,000,000 | Yes | The maximum number of rows that can be in a dataset |
| Maximum number of tags you can add to a resource | Each supported Region: 50 | No | Maximum number of tags you can add to a resource |
| Maximum number of time series per predictor | ap-south-1: 1,000,000<br><br>Each of the other supported Regions: 5,000,000 | Yes | The maximum number of time series allowed for training a predictor (number of items * number of unique values across forecast dimensions in the target time series dataset) |
| Maximum parallel running CreateAutoPredictor tasks | Each supported Region: 3 | No | The maximum number of parallel running CreateAutoPredictor tasks |
| Maximum parallel running CreateDatasetImportJob tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreateDatasetImportJob tasks |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum parallel running CreateExplainability tasks | Each supported Region: 3 | No | Maximum number of parallel running CreateExplainability tasks |
| Maximum parallel running CreateExplainabilityExport tasks | Each supported Region: 3 | No | Maximum number of parallel running CreateExplainabilityExport tasks |
| Maximum parallel running CreateForecast tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreateForecast tasks |
| Maximum parallel running CreateForecastExportJob tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreateForecastExportJob tasks |
| Maximum parallel running CreatePredictor tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreatePredictor tasks |
| Maximum parallel running CreatePredictor tasks using AutoML | Each supported Region: 3 | Yes | The maximum number of parallel running CreatePredictor tasks using AutoML |
| Maximum parallel running CreatePredictorBacktestExportJob tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreatePredictorBacktestExportJob tasks |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum parallel running CreateWhatIfAnalysis tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreateWhatIfAnalysis tasks |
| Maximum parallel running CreateWhatIfForecast tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreateWhatIfForecast tasks |
| Maximum parallel running CreateWhatIfForecastExport tasks | Each supported Region: 3 | Yes | The maximum number of parallel running CreateWhatIfForecastExport tasks |
| Maximum parallel running QueryForecast API tasks | Each supported Region: 10 | No | 10 concurrent forecasts, including 5 created with large datasets (anything over 20GB or 100,000 items). |
| Maximum parallel running Stop jobs per resource type | Each supported Region: 3 | No | Maximum number of parallel Stop jobs in progress |
| Maximum time for which a forecast can be queried on console or QueryForecast API | Each supported Region: 30 | No | Maximum time (in days) for which a forecast can be queried on console or QueryForecast API |
| The maximum number of AutoPredictors | Each supported Region: 500 | No | The maximum number of AutoPredictors that you can have in your Amazon Forecast account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| The maximum number of What-if Analyses | Each supported Region: 500 | Yes | The maximum number of What-if Analyses that you can have in your Amazon Forecast account |
| The maximum number of What-if Forecast Exports | Each supported Region: 1,000 | Yes | The maximum number of What-if Forecast Exports that you can have in your Amazon Forecast account |
| The maximum number of What-if Forecasts | Each supported Region: 100 | Yes | The maximum number of What-if Forecasts that you can have in your Amazon Forecast account |
| The maximum number of What-if Forecasts in a CreateWhatIfForecastExport task | Each supported Region: 3 | No | The maximum number of What-if Forecasts in a CreateWhatIfForecastExport task |

# Amazon Fraud Detector endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | frauddetector.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | frauddetector.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | frauddetector.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | frauddetector.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | frauddetector.ap-southeast-2.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | frauddetector.eu-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent training jobs per model | Each supported Region: 1 | No | Maximum number of concurrent training jobs per model. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Deployed model versions | Each supported Region: 5 | No | Maximum number of deployed model versions per account. |
| Detectors per account | Each supported Region: 100 | No | Maximum number of detectors per account. |
| EntityType per account | Each supported Region: 100 | No | Maximum number of EntityType per account |
| EventType per account | Each supported Region: 100 | No | Maximum number of EventType per account |
| Labels per account | Each supported Region: 100 | No | Maximum number of labels per account |
| Lists per account | Each supported Region: 100 | No | Maximum number of lists per account. |
| Models including external models per detector version | Each supported Region: 10 | No | Maximum number of models including external models per detector version. |
| Models per account | Each supported Region: 50 | No | Maximum number of models per account. |
| Outcomes per account | Each supported Region: 5,000 | No | Maximum number of outcomes per account. |
| Rate of GetPrediction requests | Each supported Region: 200 | Yes | Maximum number of GetPrediction API calls that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rules per account | Each supported Region: 5,000 | No | Maximum number of rules per account. |
| Size of GetPrediction requests | Each supported Region: 256 Kilobytes | No | Maximum size of payload per GetPrediction API call. |
| Total concurrent Event Type statistics update operations | Each supported Region: 1 | Yes | Maximum total number of concurrent Event Type statistics update operations per account. |
| Total concurrent training jobs | Each supported Region: 3 | No | Maximum total number of concurrent training jobs per account. |
| Training data size | Each supported Region: 5 Gigabytes | No | Maximum size of Fraud Detector model training data. |
| Variables per account | Each supported Region: 5,000 | No | Maximum number of variables per account. |
| Versions per detector | Each supported Region: 100 | No | Maximum number of draft versions per detector. |
| Versions per model | Each supported Region: 200 | No | Maximum number of versions per model. |

For more information, see Quotas in the *Amazon Fraud Detector User Guide*.

# FreeRTOS endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The following tables provide a list of Region-specific endpoints that FreeRTOS supports for Over-the-Air functionality. The FreeRTOS console is also supported in these Regions.

### FreeRTOS OTA Control Plane

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | iot.us-east-2.amazonaws.com<br>iot-fips.us-east-2.amazonaws.com | https<br>https | |
| US East (N. Virginia) | us-east-1 | iot.us-east-1.amazonaws.com<br>iot-fips.us-east-1.amazonaws.com | https<br>https | |
| US West (N. California) | us-west-1 | iot.us-west-1.amazonaws.com<br>iot-fips.us-west-1.amazonaws.com | https<br>https | |
| US West (Oregon) | us-west-2 | iot.us-west-2.amazonaws.com<br>iot-fips.us-west-2.amazonaws.com | https<br>https | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | iot.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | iot.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | iot.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | iot.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | iot.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | iot.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | iot.ca-central-1.amazonaws.com<br>iot-fips.ca-central-1.amazonaws.com | https<br>https | |
| Europe (Frankfurt) | eu-central-1 | iot.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | iot.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | iot.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | iot.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | iot.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | iot.me-south-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | iot.sa-east-1.amazonaws.com | HTTPS | |

## FreeRTOS OTA Data Plane

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | *prefix*.iot.us-east-2.amazonaws.com | MQTT |
| US East (N. Virginia) | us-east-1 | *prefix*.iot.us-east-1.amazonaws.com | MQTT |
| US West (N. California) | us-west-1 | *prefix*.iot.us-west-1.amazonaws.com | MQTT |
| US West (Oregon) | us-west-2 | *prefix*.iot.us-west-2.amazonaws.com | MQTT |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | *prefix*.iot.ap-east-1.amazonaws.com | MQTT |
| Asia Pacific (Mumbai) | ap-south-1 | *prefix*.iot.ap-south-1.amazonaws.com | MQTT |
| Asia Pacific (Seoul) | ap-northeast-2 | *prefix*.iot.ap-northeast-2.amazonaws.com | MQTT |
| Asia Pacific (Singapore) | ap-southeast-1 | *prefix*.iot.ap-southeast-1.amazonaws.com | MQTT |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*.iot.ap-southeast-2.amazonaws.com | MQTT |
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*.iot.ap-northeast-1.amazonaws.com | MQTT |
| Canada (Central) | ca-central-1 | *prefix*.iot.ca-central-1.amazonaws.com | MQTT |
| Europe (Frankfurt) | eu-central-1 | *prefix*.iot.eu-central-1.amazonaws.com | MQTT |
| Europe (Ireland) | eu-west-1 | *prefix*.iot.eu-west-1.amazonaws.com | MQTT |
| Europe (London) | eu-west-2 | *prefix*.iot.eu-west-2.amazonaws.com | MQTT |
| Europe (Paris) | eu-west-3 | *prefix*.iot.eu-west-3.amazonaws.com | MQTT |
| Europe (Stockholm) | eu-north-1 | *prefix*.iot.eu-north-1.amazonaws.com | MQTT |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | *prefix*.iot.me-south-1.amazonaws.com | MQTT |
| South America (São Paulo) | sa-east-1 | *prefix*.iot.sa-east-1.amazonaws.com | MQTT |

## Service quotas

### FreeRTOS OTA Resource Quotas

| Resource | Default |
|---|---|
| File size | 16MB |

### FreeRTOS OTA Throttling

| API | Transactions Per Second |
|---|---|
| CreateOTAUpdate | 10 TPS |
| DeleteOTAUpdate | 5 TPS |
| GetOTAUpdate | 15 TPS |
| ListOTAUpdates | 15 TPS |

# Amazon FSx endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | fsx.us-east-2.amazonaws.com | HTTPS |
| | | fsx-fips.us-east-2.api.aws | HTTPS |
| | | fsx-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | fsx.us-east-1.amazonaws.com | HTTPS |
| | | fsx-fips.us-east-1.api.aws | HTTPS |
| | | fsx-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | fsx.us-west-1.amazonaws.com | HTTPS |
| | | fsx-fips.us-west-1.api.aws | HTTPS |
| | | fsx-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | fsx.us-west-2.amazonaws.com | HTTPS |
| | | fsx-fips.us-west-2.api.aws | HTTPS |
| | | fsx-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | fsx.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-east-1 | fsx.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | fsx.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | fsx.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | fsx.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | fsx.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | fsx.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | fsx.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | fsx.ap-northeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | fsx.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | fsx.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | fsx.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | fsx.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | fsx.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | fsx.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | fsx.ca-central-1.amazonaws.com<br>fsx-fips.ca-central-1.amazonaws.com<br>fsx-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | fsx.ca-west-1.amazonaws.com<br>fsx-fips.ca-west-1.amazonaws.com<br>fsx-fips.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | fsx.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | fsx.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | fsx.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | fsx.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | fsx.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | fsx.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | fsx.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | fsx.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | fsx.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | fsx.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | fsx.me-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | fsx.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | fsx.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | fsx.us-gov-east-1.amazonaws.com | HTTPS |
| | | fsx-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | fsx-fips.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | fsx.us-gov-west-1.amazonaws.com | HTTPS |
| | | fsx-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | fsx-fips.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Lustre Cache_1 caches | Each supported Region: 100 | Yes | The maximum number of Amazon File Cache caches with cache type Lustre and deployment type Cache_1 that you can create in this account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Lustre Cache_1 storage capacity | Each supported Region: 100,800 | Yes | The maximum amount of storage capacity (in GiB) that you can configure in this account for all Amazon File Cache caches with cache type Lustre and deployment type Cache_1. |
| Lustre Persistent HDD storage capacity (per file system) | Each supported Region: 102,000 | Yes | The maximum amount of HDD storage capacity (in GiB) that you can configure for an Amazon FSx for Lustre persistent file system. |
| Lustre Persistent Intelligent-Tiering SSD read cache storage capacity | Each supported Region: 100,800 | Yes | The maximum amount of provisioned SSD read cache storage capacity (in GiB) that you can configure for all Amazon FSx for Lustre Intelligent-Tiering file systems in this account. |
| Lustre Persistent Intelligent-Tiering throughput capacity | Each supported Region: 100,000 | Yes | The total amount of throughput capacity (in MBps) allowed for all Amazon FSx for Lustre Intelligent-Tiering file systems in this account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Lustre Persistent_1 file systems | Each supported Region: 100 | Yes | The maximum number of Amazon FSx for Lustre persistent_1 file systems that you can create in this account |
| Lustre Persistent_1 storage capacity | Each supported Region: 100,800 | Yes | The maximum amount of storage capacity (in GiB) that you can configure for all Amazon FSx for Lustre persistent_1 file systems in this account. |
| Lustre Persistent_2 file systems | Each supported Region: 100 | Yes | The maximum number of Amazon FSx for Lustre persistent_2 file systems that you can create in this account |
| Lustre Persistent_2 storage capacity | Each supported Region: 100,800 | Yes | The maximum amount of storage capacity (in GiB) that you can configure for all Amazon FSx for Lustre persistent_2 file systems in this account. |
| Lustre Scratch file systems | Each supported Region: 100 | Yes | The maximum number of Amazon FSx for Lustre scratch file systems that you can create in this account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Lustre Scratch storage capacity | Each supported Region: 100,800 | Yes | The maximum amount of storage capacity (in GiB) that you can configure for all Amazon FSx for Lustre scratch file systems in this account. |
| Lustre backups | Each supported Region: 500 | Yes | The maximum number of user-initiated backups that you can have for all Amazon FSx for Lustre file systems in this account. |
| ONTAP SSD IOPS | Each supported Region: 2,000,000 | Yes | The total amount of SSD IOPS allowed for all Amazon FSx for NetApp ONTAP file systems in this account. |
| ONTAP SSD storage capacity | Each supported Region: 1,048,576 | Yes | The maximum amount of SSD storage capacity (in GiB) for all Amazon FSx for NetApp ONTAP file systems that you can have in this account. |
| ONTAP backups | Each supported Region: 10,000 | Yes | The maximum number of user-initiated backups for all Amazon FSx for NetApp ONTAP file systems that you can have in this account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ONTAP file systems | Each supported Region: 100 | Yes | The maximum number of Amazon FSx for NetApp ONTAP file systems that you can create in this account. |
| ONTAP throughput capacity | Each supported Region: 10,240 | Yes | The total amount of throughput capacity (in MBps) allowed for all Amazon FSx for NetApp ONTAP file systems in this account. |
| OpenZFS SSD storage capacity | us-east-1: 262,144<br><br>us-east-2: 262,144<br><br>us-west-2: 262,144<br><br>Each of the other supported Regions: 65,536 | Yes | The maximum amount of SSD storage capacity (in GiB) that you can configure for all Amazon FSx for OpenZFS file systems in this account. |
| OpenZFS SSD storage capacity (per file system) | Each supported Region: 524,288 | Yes | The maximum amount of SSD storage capacity (in GiB) that you can configure for an Amazon FSx for OpenZFS file system. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| OpenZFS backups | Each supported Region: 10,000 | Yes | The maximum number of user-initiated backups for all Amazon FSx for OpenZFS file systems that you can have in this account. |
| OpenZFS disk IOPS | Each supported Region: 400,000 | Yes | The total amount of disk IOPS allowed for all Amazon FSx for OpenZFS file systems in this account. |
| OpenZFS file systems | Each supported Region: 100 | Yes | The maximum number of Amazon FSx for OpenZFS file systems that you can create in this account. |
| OpenZFS provisioned SSD read cache storage capacity | us-east-1: 262,144<br>us-east-2: 262,144<br>us-west-2: 262,144<br>Each of the other supported Regions: 65,536 | Yes | The maximum amount of provisioned SSD read cache storage capacity (in GiB) that you can configure for all Amazon FSx for OpenZFS file systems in this account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| OpenZFS throughput capacity | Each supported Region: 10,240 | Yes | The total amount of throughput capacity (in MBps) allowed for all Amazon FSx for OpenZFS file systems in this account. |
| Windows HDD storage capacity | Each supported Region: 524,288 | Yes | The maximum amount of HDD storage capacity (in GiB) allowed for all Amazon FSx for Windows File Server file systems in this account. |
| Windows SSD storage capacity | Each supported Region: 524,288 | Yes | The maximum amount of SSD storage capacity (in GiB) for all Amazon FSx for Windows File Server file systems that you can have in this account. |
| Windows backups | Each supported Region: 500 | Yes | The maximum number of user-initiated backups for all Amazon FSx for Windows File Server file systems that you can have in this account. |
| Windows file systems | Each supported Region: 100 | Yes | The maximum number of Amazon FSx for Windows Server file systems that you can create in this account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Windows throughput capacity | Each supported Region: 10,240 | Yes | The total amount of throughput capacity (in MBps) allowed for all Amazon FSx for Windows file systems in this account. |
| Windows total SSD IOPS | Each supported Region: 500,000 | Yes | The total amount of SSD IOPS allowed for all Amazon FSx for Windows File Server file systems that you can have in this account. |

For more information, see the following:

- FSx for Lustre quotas in the *Amazon FSx for Lustre User Guide*
- FSx for ONTAP quotas in the *FSx for ONTAP User Guide*
- FSx for OpenZFS quotas in the *FSx for OpenZFS User Guide*
- FSx for Windows quotas in the *Amazon FSx for Windows File Server User Guide*

## Amazon GameLift Servers endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Note**

The following service endpoints table lists the AWS Regions that are available as Amazon GameLift Servers fleet home Regions. For more on Region availability for Amazon GameLift Servers features, including multi-location fleets, remote locations, and Local Zones, see Amazon GameLift Servers service locations.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | gamelift.us-east-2.amazonaws.com | HTTPS |
| | | gamelift.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | gamelift.us-east-1.amazonaws.com | HTTPS |
| | | gamelift.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | gamelift.us-west-1.amazonaws.com | HTTPS |
| | | gamelift.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | gamelift.us-west-2.amazonaws.com | HTTPS |
| | | gamelift.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | gamelift.ap-south-1.amazonaws.com | HTTPS |
| | | gamelift.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | gamelift.ap-northeast-2.amazonaws.com | HTTPS |
| | | gamelift.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | gamelift.ap-southeast-1.amazonaws.com<br>gamelift.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | gamelift.ap-southeast-2.amazonaws.com<br>gamelift.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | gamelift.ap-northeast-1.amazonaws.com<br>gamelift.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | gamelift.ca-central-1.amazonaws.com<br>gamelift.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | gamelift.eu-central-1.amazonaws.com<br>gamelift.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | gamelift.eu-west-1.amazonaws.com<br>gamelift.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | gamelift.eu-west-2.amazonaws.com<br>gamelift.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | gamelift.sa-east-1.amazonaws.com<br>gamelift.sa-east-1.api.aws | HTTPS<br>HTTPS | |

# UDP ping beacon endpoints

Amazon GameLift Servers UDP ping beacons allow you to measure network latency between your players and AWS locations. You can use these UDP ping beacon endpoints to collect latency data and make informed decisions about which hosting locations provide the best gaming experience for your players. The following table lists the UDP ping beacon endpoints available for each AWS location and Local Zone that Amazon GameLift Servers supports.

### UDP ping beacons by location

| Location name | Location code | Endpoint | Protocol | IPv6 Support |
|---|---|---|---|---|
| **North America** | | | | |
| US East (N. Virginia) | us-east-1 | gamelift-ping.us-east-1.api.aws:7770 | UDP | Yes |
| US East (Ohio) | us-east-2 | gamelift-ping.us-east-2.api.aws:7770 | UDP | Yes |
| US West (N. California) | us-west-1 | gamelift-ping.us-west-1.api.aws:7770 | UDP | Yes |
| US West (Oregon) | us-west-2 | gamelift-ping.us-west-2.api.aws:7770 | UDP | Yes |
| Canada (Central) | ca-central-1 | gamelift-ping.ca-central-1.api.aws:7770 | UDP | Yes |
| **US Local Zones** | | | | |
| US West (Los Angeles) | us-west-2-lax-1 | gamelift-ping-lax.us-west-2.api.aws:7770 | UDP | Yes |
| US East (Chicago) | us-east-1-chi-1 | gamelift-ping-chi.us-east-1.api.aws:7770 | UDP | No |

| Location name | Location code | Endpoint | Protocol | IPv6 Support |
|---|---|---|---|---|
| US East (Houston) | us-east-1-iah-1 | gamelift-ping-iah.us-east-1.api.aws:7770 | UDP | No |
| US East (Dallas) | us-east-1-dfw-1 | gamelift-ping-dfw.us-east-1.api.aws:7770 | UDP | No |
| US West (Denver) | us-west-2-den-1 | gamelift-ping-den.us-west-2.api.aws:7770 | UDP | No |
| US East (Atlanta) | us-east-1-atl-1 | gamelift-ping-atl.us-east-1.api.aws:7770 | UDP | No |
| US West (Phoenix) | us-west-2-phx-1 | gamelift-ping-phx.us-west-2.api.aws:7770 | UDP | No |
| US East (Kansas City) | us-east-1-mci-1 | gamelift-ping-mci.us-east-1.api.aws:7770 | UDP | No |
| **South America** | | | | |
| South America (São Paulo) | sa-east-1 | gamelift-ping.sa-east-1.api.aws:7770 | UDP | Yes |
| **Europe** | | | | |
| Europe (Ireland) | eu-west-1 | gamelift-ping.eu-west-1.api.aws:7770 | UDP | Yes |
| Europe (London) | eu-west-2 | gamelift-ping.eu-west-2.api.aws:7770 | UDP | Yes |
| Europe (Paris) | eu-west-3 | gamelift-ping.eu-west-3.api.aws:7770 | UDP | Yes |

| Location name | Location code | Endpoint | Protocol | IPv6 Support |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | gamelift-ping.eu-central-1.api.aws:7770 | UDP | Yes |
| Europe (Milan) | eu-south-1 | gamelift-ping.eu-south-1.api.aws:7770 | UDP | Yes |
| Europe (Stockholm) | eu-north-1 | gamelift-ping.eu-north-1.api.aws:7770 | UDP | Yes |
| **Asia Pacific** | | | | |
| Asia Pacific (Malaysia) | ap-southeast-5 | gamelift-ping.ap-southeast-5.api.aws:7770 | UDP | Yes |
| Asia Pacific (Mumbai) | ap-south-1 | gamelift-ping.ap-south-1.api.aws:7770 | UDP | Yes |
| Asia Pacific (Hong Kong) | ap-east-1 | gamelift-ping.ap-east-1.api.aws:7770 | UDP | Yes |
| Asia Pacific (Thailand) | ap-southeast-7 | gamelift-ping.ap-southeast-7.api.aws:7770 | UDP | Yes |
| Asia Pacific (Osaka) | ap-northeast-3 | gamelift-ping.ap-northeast-3.api.aws:7770 | UDP | Yes |
| Asia Pacific (Seoul) | ap-northeast-2 | gamelift-ping.ap-northeast-2.api.aws:7770 | UDP | Yes |
| Asia Pacific (Singapore) | ap-southeast-1 | gamelift-ping.ap-southeast-1.api.aws:7770 | UDP | Yes |
| Asia Pacific (Sydney) | ap-southeast-2 | gamelift-ping.ap-southeast-2.api.aws:7770 | UDP | Yes |
| Asia Pacific (Tokyo) | ap-northeast-1 | gamelift-ping.ap-northeast-1.api.aws:7770 | UDP | Yes |

| Location name | Location code | Endpoint | Protocol | IPv6 Support |
|---|---|---|---|---|
| **Middle East** | | | | |
| Middle East (Bahrain) | me-south-1 | gamelift-ping.me-south-1.api.aws:7770 | UDP | Yes |
| **Africa** | | | | |
| Africa (Cape Town) | af-south-1 | gamelift-ping.af-south-1.api.aws:7770 | UDP | Yes |
| Africa (Lagos) | af-south-1-los-1 | gamelift-ping-los.af-south-1.api.aws:7770 | UDP | No |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Aliases per region | Each supported Region: 100 | Yes | The maximum number of aliases allowed per region. |
| Anywhere fleets per region | Each supported Region: 30 | Yes | The maximum number of Anywhere fleets allowed (in any status) per region. |
| Build capacity | Each supported Region: 100 Gigabytes | No | The maximum capacity (in gigabytes) available for all uploaded game builds combined per region. To free storage space, delete unused builds as needed. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Builds per region | Each supported Region: 1,000 | Yes | The maximum number of game server builds allowed (in any status) per region. |
| Compute per Anywhere fleet | Each supported Region: 100 | Yes | The maximum number of registered compute resources allowed per Anywhere fleet. |
| Custom locations per region | Each supported Region: 20 | Yes | The maximum number of custom locations allowed per region. |
| Game server groups per region | Each supported Region: 20 | Yes | The maximum number of game server groups allowed per region. |
| Game servers per game server group | Each supported Region: 1,000 | Yes | The maximum number of game servers allowed per game server group. |
| Game session log file size | Each supported Region: 200 Megabytes | No | The maximum file size (in megabytes) allowed for game session logs that are uploaded to Amazon GameLift Servers at the conclusion of a game session. |
| Game session queues per region | Each supported Region: 100 | Yes | The maximum number of game session queues allowed per region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Key-value pairs per string to double map matchmaking player attribute | Each supported Region: 10 | No | The maximum number of key-value pairs in a string to double map (SDM) matchmaking player attribute. |
| Locations in a fleet per region | Each supported Region: 10 | Yes | The maximum number of locations allowed (in any status) in a fleet per region. |
| Managed EC2 fleet EBS volume | Each supported Region: 50 Gigabytes | Yes | The provisioned EBS volume capacity (in gigabytes) for managed EC2 fleet instances. |
| Managed EC2 fleets per region | Each supported Region: 20 | Yes | The maximum number of managed EC2 fleets allowed (in any status) per region. |
| Matchmaking configurations per region | Each supported Region: 300 | Yes | The maximum number of matchmaking configurations allowed per region. |
| Matchmaking rule sets per region | Each supported Region: 1,000 | Yes | The maximum number of matchmaking rule sets allowed per region. |
| Maximum NewGameSessionsPerCreator per fleet configuration | Each supported Region: 10 | Yes | The maximum new game sessions per creator allowed in a fleets resource policy configuration. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum PolicyPeriodInMinutes per fleet configuration | Each supported Region: 60 | Yes | The maximum period (in minutes) allowed in a fleets resource policy configuration. |
| Player attributes per matchmaking player | Each supported Region: 10 | No | The maximum number of player attributes for each player in a matchmaking ticket. |
| Player session timeout | Each supported Region: 60 Seconds | Yes | The maximum time (in seconds) for a reserved player session to timeout pending a valid player connection. |
| Player sessions per game session | Each supported Region: 200 | No | The maximum number of player sessions that can join a game session. |
| Players per matchmaking ticket | Each supported Region: 10 | No | The maximum number of players that can be included in a matchmaking ticket. |
| Queue destinations per game session queue | Each supported Region: 10 | Yes | The maximum number of queue destinations allowed per game session queue. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Script capacity | Each supported Region: 100 Gigabytes | No | The maximum capacity (in gigabytes) available for all uploaded scripts combined per region. To free storage space, delete unused builds as needed. |
| Scripts per region | Each supported Region: 1,000 | Yes | The maximum number of game server scripts allowed per region. |
| Server processes per instance (Server SDK v3 and up) | Each supported Region: 50 | No | The maximum number of concurrent server processes that can run on a single instance when using the Server SDK for Amazon GameLift Servers version 3 or later. |
| Strings per string list matchmaking player attribute | Each supported Region: 100 | No | The maximum number of strings in a string list (SL) matchmaking player attribute. |

## Amazon GameLift Streams endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

> **Note**
>
> The following service endpoints table lists the AWS Regions that are available as Amazon GameLift Streams primary locations. Additional Regions are available as remote locations. For more on Region availability for Amazon GameLift Streams features, including multi-location stream groups and remote locations, see Regions, quotas, and limitations in the *Amazon GameLift Streams Developer Guide*.

## Service endpoints

Amazon GameLift Streams supports dual-stack service endpoints, allowing clients and resources to interact with the service using IPv6 or IPv4.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | gameliftstreams.us-east-2.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | gameliftstreams.us-west-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | gameliftstreams.ap-northeast-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | gameliftstreams.eu-central-1.api.aws | HTTPS |

## Service quotas

In the following table, the GPU quotas are all 0 by default. However, your account's applied quotas might be different. To check, sign in to the AWS Management Console and open the Service Quotas console to Amazon GameLift Streams, where you can review your current quotas in the **Applied**

**account-level quota value** column. You can also use the Service Quotas console to request an increase to a quota.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Application size (GiB) | Each supported Region: 100 | Yes | The maximum total size (in GiB) of an application, in this account. Note that a gibibyte (GiB) equals 1024*1024*1024 bytes. |
| Applications | Each supported Region: 20 | Yes | The maximum number of applications that you can create in this account, per AWS Region. |
| Files per application | Each supported Region: 30,000 | Yes | The maximum number of files that you can have in an application, in this account. |
| Gen4n GPUs, ap-northeast-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, ap-northeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the ap-northeast-2 |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, ap-south-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the ap-south-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, ap-southeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the ap-southeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen4n GPUs, eu-central-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, eu-north-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, eu-west-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the eu-west-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen4n GPUs, eu-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the eu-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, sa-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the sa-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, us-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen4n GPUs, us-east-2 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen4n GPUs, us-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen4n GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen4n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, ap-northeast-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen5n GPUs, ap-northeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the ap-northeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, ap-south-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the ap-south-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen5n GPUs, ap-southeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the ap-southeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, eu-central-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen5n GPUs, eu-north-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, eu-west-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the eu-west-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, eu-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the eu-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen5n GPUs, sa-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the sa-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, us-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen5n GPUs, us-east-2 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen5n GPUs, us-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen5n GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen5n_high", support streaming more than one session per GPU. |
| Gen6e GPUs, ap-northeast-1 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |
| Gen6e GPUs, ap-northeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the ap-northeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen6e GPUs, eu-central-1 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |
| Gen6e GPUs, eu-north-1 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |
| Gen6e GPUs, us-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen6e GPUs, us-east-2 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |
| Gen6e GPUs, us-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen6e GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6e_high", support streaming more than one session per GPU. |
| Gen6n GPUs, ap-northeast-1 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the ap-northeast-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen6n GPUs, ap-northeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the ap-northeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |
| Gen6n GPUs, ap-south-1 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the ap-south-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen6n GPUs, ap-southeast-2 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the ap-southeast-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |
| Gen6n GPUs, eu-central-1 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the eu-central-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen6n GPUs, eu-north-1 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the eu-north-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |
| Gen6n GPUs, eu-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the eu-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |
| Gen6n GPUs, sa-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the sa-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Gen6n GPUs, us-east-1 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the us-east-1 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |
| Gen6n GPUs, us-east-2 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the us-east-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |
| Gen6n GPUs, us-west-2 | Each supported Region: 0 | Yes | The maximum number of Gen6n GPUs you can configure for streaming in the us-west-2 location across all stream groups in this account. Multi-tenant stream classes, such as "Gen6n_high", support streaming more than one session per GPU. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Stream groups | Each supported Region: 5 | Yes | The maximum number of stream groups that you can create in this account, per AWS Region. A stream group is a collection of compute resources that stream your application to end users. |

# AWS Global Accelerator endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoint (global)

| Region Name | Region | Endpoint | Protocol | Amazon Route 53 Hosted Zone ID* |
|---|---|---|---|---|
| US West (Oregon) Region | us-west-2 | globalaccelerator.amazonaws.com | HTTPS | Z2BJ6XQ5FK7U4H |

| Region Name | Region | Endpoint | Protocol | Amazon Route 53 Hosted Zone ID* |
|---|---|---|---|---|
| US West (Oregon) Region | us-west-2 | globalaccelerator-fips.us-west-2.amazonaws.com | HTTPS | Z01882459 3QZ67JH63 2G |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Cross-account attachments per AWS account | Each supported Region: 10 | Yes | The maximum number of cross-account attachments for each AWS account. |
| Custom routing accelerators per AWS account | Each supported Region: 10 | Yes | The maximum number of custom routing accelerators for each AWS account. |
| Endpoint groups per accelerator | Each supported Region: 42 | No | The maximum number of endpoint groups per accelerator. |
| Endpoint resources per cross-account attachment | Each supported Region: 500 | No | The maximum number of endpoint resources per cross-account attachment. |
| Endpoints per endpoint group - Application Load Balancers | Each supported Region: 10 | No | The maximum number of Application Load Balancers in an endpoint |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | group containing only ALB endpoints. |
| Endpoints per endpoint group - EC2 instances | Each supported Region: 10 | Yes | The maximum number of EC2 instances in an endpoint group containing only EC2 instance endpoints. |
| Endpoints per endpoint group - Elastic IP addresses | Each supported Region: 10 | Yes | The maximum number of Elastic IP addresses in an endpoint group containing only Elastic IP address endpoints. |
| Endpoints per endpoint group - Network Load Balancers | Each supported Region: 10 | Yes | The maximum number of Network Load Balancers in an endpoint group containing only NLB endpoints. |
| Endpoints per endpoint group - VPC subnets | Each supported Region: 10 | Yes | The maximum number of VPC subnets in an endpoint group containing only subnet endpoints. |
| Endpoints per endpoint group - more than one endpoint type | Each supported Region: 10 | No | The maximum number of endpoints in an endpoint group containing more than one endpoint type. |
| Listeners per accelerator | Each supported Region: 10 | Yes | The maximum number of listeners for each accelerator. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Port overrides per endpoint group | Each supported Region: 10 | Yes | The maximum number of port overrides for each endpoint group. |
| Port ranges per listener | Each supported Region: 10 | No | The maximum number of port ranges for each listener. |
| Principals per cross-account attachment | Each supported Region: 10 | Yes | The maximum number of principals per cross-account attachment. |
| Standard accelerators per AWS account | Each supported Region: 20 | Yes | The maximum number of standard accelerators for each AWS account. |
| Tags per accelerator | Each supported Region: 50 | No | The maximum number of tags for each accelerator. |

# AWS Global Networks for Transit Gateways Quotas (Network Manager)

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Africa (Cape Town) | af-south-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | networkmanager.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS | |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS | |
| | | networkmanager.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Milan) | eu-south-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | networkmanager.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager-fips.us-west-2.api.aws | HTTPS |
| | | networkmanager-fips.us-west-2.amazonaws.com | HTTPS |
| | | networkmanager.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | networkmanager.us-west-2.amazonaws.com<br>networkmanager-fips.us-west-2.api.aws<br>networkmanager-fips.us-west-2.amazonaws.com<br>networkmanager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | networkmanager.us-gov-west-1.amazonaws.com<br>networkmanager.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | networkmanager.us-gov-west-1.amazonaws.com<br>networkmanager.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

For a list of the quotas, see AWS Global Networks for Transit Gateways Quotas in the *AWS Global Networks for Transit Gateways User Guide*.

# AWS Glue endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | glue.us-east-2.amazonaws.com | HTTPS | |
| | | glue-fips.us-east-2.api.aws | HTTPS | |
| | | glue-fips.us-east-2.amazonaws.com | HTTPS | |
| | | glue.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | glue.us-east-1.amazonaws.com | HTTPS | |
| | | glue-fips.us-east-1.api.aws | HTTPS | |
| | | glue-fips.us-east-1.amazonaws.com | HTTPS | |
| | | glue.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | glue.us-west-1.amazonaws.com | HTTPS | |
| | | glue-fips.us-west-1.api.aws | HTTPS | |
| | | glue-fips.us-west-1.amazonaws.com | HTTPS | |
| | | glue.us-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | glue.us-west-2.amazonaws.com | HTTPS |
| | | glue-fips.us-west-2.api.aws | HTTPS |
| | | glue-fips.us-west-2.amazonaws.com | HTTPS |
| | | glue.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | glue.af-south-1.amazonaws.com | HTTPS |
| | | glue.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | glue.ap-east-1.amazonaws.com | HTTPS |
| | | glue.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | glue.ap-south-2.amazonaws.com | HTTPS |
| | | glue.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | glue.ap-southeast-3.amazonaws.com | HTTPS |
| | | glue.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | glue.ap-southeast-5.amazonaws.com | HTTPS |
| | | glue.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | glue.ap-southeast-4.amazonaws.com | HTTPS |
| | | glue.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | glue.ap-south-1.amazonaws.com<br>glue.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | glue.ap-southeast-6.amazonaws.com<br>glue.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | glue.ap-northeast-3.amazonaws.com<br>glue.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | glue.ap-northeast-2.amazonaws.com<br>glue.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | glue.ap-southeast-1.amazonaws.com<br>glue.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | glue.ap-southeast-2.amazonaws.com<br>glue.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | glue.ap-east-2.amazonaws.com<br>glue.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | glue.ap-southeast-7.amazonaws.com<br>glue.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | glue.ap-northeast-1.amazonaws.com | HTTPS |
| | | glue.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | glue.ca-central-1.amazonaws.com | HTTPS |
| | | glue-fips.ca-central-1.api.aws | HTTPS |
| | | glue-fips.ca-central-1.amazonaws.com | HTTPS |
| | | glue.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | glue.ca-west-1.amazonaws.com | HTTPS |
| | | glue-fips.ca-west-1.api.aws | HTTPS |
| | | glue-fips.ca-west-1.amazonaws.com | HTTPS |
| | | glue.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | glue.eu-central-1.amazonaws.com | HTTPS |
| | | glue.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | glue.eu-west-1.amazonaws.com | HTTPS |
| | | glue.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | glue.eu-west-2.amazonaws.com | HTTPS |
| | | glue.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | glue.eu-south-1.amazonaws.com | HTTPS |
| | | glue.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | glue.eu-west-3.amazonaws.com | HTTPS |
| | | glue.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | glue.eu-south-2.amazonaws.com<br>glue.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | glue.eu-north-1.amazonaws.com<br>glue.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | glue.eu-central-2.amazonaws.com<br>glue.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | glue.il-central-1.amazonaws.com<br>glue.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | glue.mx-central-1.amazonaws.com<br>glue.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | glue.me-south-1.amazonaws.com<br>glue.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | glue.me-central-1.amazonaws.com<br>glue.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | glue.sa-east-1.amazonaws.com<br>glue.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | glue.us-gov-east-1.amazonaws.com | HTTPS |
| | | glue-fips.us-gov-east-1.api.aws | HTTPS |
| | | glue-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | glue.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | glue.us-gov-west-1.amazonaws.com | HTTPS |
| | | glue-fips.us-gov-west-1.api.aws | HTTPS |
| | | glue-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | glue.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent compute capacity for Ray worker instances in M-DPUs per account. | Each supported Region: 50 | No | The maximum number of data processing units that can be concurrently used by Ray jobs or interactive sessions in this account. The capacity provided by a Ray worker instance type is measured in memory optimized data processing units (M-DPUs). |
| Concurrent machine learning task runs per transform | Each supported Region: 3 | Yes | The maximum number of concurrent task runs |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | per machine learning transform for this account. |
| Global max tags | Each supported Region: 10,000 | Yes | Catalog: global max tags |
| Global maximum administrators | Each supported Region: 100 | Yes | Catalog: global maximum administrators |
| Global maximum paths | Each supported Region: 30,000 | Yes | Catalog: global maximum paths |
| Global s3 path max length | Each supported Region: 2,000 | Yes | Catalog: global s3 path max length |
| Label file size | Each supported Region: 10 Megabytes | Yes | The maximum file size of an individual label file that can be imported. |
| Max DPUs per dev endpoint | Each supported Region: 50 | Yes | The maximum number of DPUs in your development endpoint. |
| Max Named Tag Expressions | Each supported Region: 25,000 | Yes | The maximum number of named tag expression allowed per account |
| Max concurrent data quality ruleset evaluation runs per account | Each supported Region: 50 | Yes | The maximum number of concurrent data quality ruleset evaluation runs in your account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max concurrent data quality ruleset recommendation runs per account | Each supported Region: 25 | Yes | The maximum number of concurrent data quality ruleset recommendation runs in your account. |
| Max concurrent job runs per account | Each supported Region: 2,000 | Yes | The maximum number of concurrent job runs in your account. |
| Max concurrent job runs per job | Each supported Region: 4,000 | Yes | The maximum number of concurrent job runs for a job. |
| Max concurrent observations generation runs per account | Each supported Region: 10 | Yes | The maximum number of concurrent observations generation runs |
| Max connection per account | Each supported Region: 1,000 | Yes | The maximum number of connections in your account. |
| Max data quality rulesets per account | Each supported Region: 1,000,000 | Yes | The maximum number of data quality rulesets per account. |
| Max databases per account | Each supported Region: 10,000 | Yes | The maximum number of databases in your account. |
| Max databases per catalog | Each supported Region: 10,000 | Yes | The maximum number of databases per catalog. |
| Max development endpoint per account | Each supported Region: 25 | Yes | The maximum number of development endpoints in your account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max functions per account | Each supported Region: 100 | Yes | The maximum number of functions in your account. |
| Max functions per database | Each supported Region: 100 | Yes | The maximum number of functions per database. |
| Max jobs per account | Each supported Region: 2,000 | Yes | The maximum number of jobs in your account. |
| Max jobs per trigger | Each supported Region: 50 | Yes | The maximum number of jobs that a trigger can start. |
| Max partitions per account | Each supported Region: 20,000,000 | Yes | The maximum number of partitions in your account. |
| Max partitions per table | Each supported Region: 10,000,000 | Yes | The maximum number of partitions per table. |
| Max security configurations per account | Each supported Region: 250 | Yes | The maximum number of security configurations in your account. |
| Max spare compute capacity consumed in data processing units (DPUs) per account. | Each supported Region: 4,000 | Yes | The maximum spare compute capacity in data processing units (DPUs) you can use concurrently in your account. |
| Max table versions per account | Each supported Region: 1,000,000 | Yes | The maximum number of table versions in your account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max table versions per table | Each supported Region: 100,000 | Yes | The maximum number of table versions per table. |
| Max tables per account | Each supported Region: 1,000,000 | Yes | The maximum number of tables in your account. |
| Max tables per database | Each supported Region: 200,000 | Yes | The maximum number of tables per database. |
| Max task DPUs per account | us-east-1: 1,000<br>us-east-2: 1,000<br>us-west-2: 1,000<br>ap-northeast-1: 1,000<br>ap-southeast-2: 1,000<br>eu-west-1: 1,000<br>Each of the other supported Regions: 500 | Yes | The maximum compute capacity in data processing units (DPUs) you can use concurrently in your account. |
| Max triggers per account | Each supported Region: 1,000 | Yes | The maximum number of triggers in your account. |
| Maximum number of in-flight completion requests allowed per account | Each supported Region: 10 | Yes | The maximum number of in-flight completion requests allowed per account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of queued job runs | Each supported Region: 1,000 | Yes | Orchestration: maximum number of queued job runs |
| Number of Schema Registries. | Each supported Region: 100 | No | The maximum number of Schema Registries per AWS Region for this account. |
| Number of Schema Versions. | Each supported Region: 10,000 | No | The maximum number of Schema Versions per AWS Region for this account. |
| Number of column statistics tasks running concurrently per account | Each supported Region: 50 | Yes | The maximum number of column statistics tasks running concurrently in your account. |
| Number of connection types | Each supported Region: 1,000 | Yes | The maximum number of connection types per AWS Region for this account. |
| Number of crawlers per account | Each supported Region: 1,000 | Yes | The maximum number of crawlers in your account. |
| Number of crawlers running concurrently per account | Each supported Region: 150 | Yes | The maximum number of crawlers running concurrently in your account. |
| Number of integrations | Each supported Region: 40 | Yes | The maximum number of integrations per AWS Region for this account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of machine learning transforms | Each supported Region: 100 | Yes | The maximum number of machine learning transforms for this account. |
| Number of metadata key value pairs per Schema Version. | Each supported Region: 10 | No | The maximum number of Schema Version metadata key value pairs per Schema Version. |
| Number of workflows | Each supported Region: 1,000 | Yes | The maximum number of workflows in your account. |
| The maximum number of concurrent materialized view refresh task runs. | Each supported Region: 50 | Yes | The maximum number of concurrent materialized view refresh task runs. |
| Total concurrent machine learning task runs for transforms per account | Each supported Region: 30 | Yes | The total number of concurrent machine learning transform task runs for machine learning transforms for this account. |

For more information, see AWS Glue in the *AWS GovCloud (US) User Guide*.

## AWS Glue DataBrew endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | databrew.us-east-2.amazonaws.com<br>databrew-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | databrew.us-east-1.amazonaws.com<br>databrew-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | databrew.us-west-1.amazonaws.com<br>databrew-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | databrew.us-west-2.amazonaws.com<br>databrew-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | databrew.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | databrew.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | databrew.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | databrew.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | databrew.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | databrew.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | databrew.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | databrew.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | databrew.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | databrew.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | databrew.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | databrew.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | databrew.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | databrew.eu-north-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | databrew.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | databrew.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent jobs per AWS account | Each supported Region: 10 | Yes | The maximum number of jobs that you can run at the same time in this AWS account. |
| Datasets per AWS account | Each supported Region: 100 | Yes | The maximum number of datasets that you can create in this AWS account. |
| Jobs per AWS account | Each supported Region: 100 | Yes | The maximum number of jobs that you can create in this AWS account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Node capacity per AWS account | Each supported Region: 300 | Yes | The maximum number of nodes available to jobs running in this AWS account. |
| Open projects per AWS account | Each supported Region: 10 | Yes | The maximum number of projects that you can open concurrently in this AWS account. |
| Projects per AWS account | Each supported Region: 100 | Yes | The maximum number of projects that you can create in this AWS account. |
| Recipes per AWS account | Each supported Region: 500 | Yes | The maximum number of recipes that you can create in this AWS account. |
| Rules per ruleset | Each supported Region: 100 | Yes | The maximum number of rules that you can have in a ruleset. |
| Rulesets per AWS account | Each supported Region: 100 | Yes | The maximum number of rulesets that you can create in this AWS account. |
| Rulesets per dataset | Each supported Region: 10 | Yes | The maximum number of rulesets that you can create for a dataset. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Schedules per AWS account | Each supported Region: 10 | Yes | The maximum number of schedules that you can create in this AWS account. |
| Versions per recipe | Each supported Region: 100 | Yes | The maximum number of versions that you can create for a recipe. |

# Amazon Managed Grafana endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | grafana.us-east-2.amazonaws.com | HTTPS |
| | | grafana.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | grafana.us-east-1.amazonaws.com | HTTPS |
| | | grafana.us-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | grafana.us-west-2.amazonaws.com | HTTPS | |
| | | grafana.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | grafana.ap-northeast-2.amazonaws.com | HTTPS | |
| | | grafana.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | grafana.ap-southeast-1.amazonaws.com | HTTPS | |
| | | grafana.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | grafana.ap-southeast-2.amazonaws.com | HTTPS | |
| | | grafana.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | grafana.ap-northeast-1.amazonaws.com | HTTPS | |
| | | grafana.ap-northeast-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | grafana.eu-central-1.amazonaws.com | HTTPS | |
| | | grafana.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | grafana.eu-west-1.amazonaws.com | HTTPS | |
| | | grafana.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | grafana.eu-west-2.amazonaws.com | HTTPS | |
| | | grafana.eu-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | grafana.us-gov-east-1.amazonaws.com | HTTPS | |
| | | grafana-fips.us-gov-east-1.api.aws | HTTPS | |
| | | grafana-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | grafana.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | grafana.us-gov-west-1.amazonaws.com | HTTPS | |
| | | grafana-fips.us-gov-west-1.api.aws | HTTPS | |
| | | grafana-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | grafana.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

Amazon Managed Grafana has the following quotas.

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Number of workspaces | Each supported Region: 5 | Yes | The maximum number of workspaces that you can have in this account in the current region. |
| Rate of AssociateLicense requests | Each supported Region: 1 per second | No | The maximum number of AssociateLicense requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateWorkspace requests | Each supported Region: 1 per second | No | The maximum number of CreateWorkspace requests that you can make, per second, in this account in the current region. |
| Rate of DeleteWorkspace requests | Each supported Region: 1 per second | No | The maximum number of DeleteWorkspace requests that you can make, per second, in this account in the current region. |
| Rate of DescribeWorkspace requests | Each supported Region: 5 per second | No | The maximum number of DescribeWorkspace requests that you can make, per second, in this account in the current region. |
| Rate of DescribeWorkspaceAuthentication requests | Each supported Region: 1 per second | No | The maximum number of DescribeWorkspaceAuthentication requests that you can make, per second, in this account in the current region. |
| Rate of DisassociateLicense requests | Each supported Region: 1 per second | No | The maximum number of DisassociateLicense requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Rate of ListPermissions requests | Each supported Region: 10 per second | No | The maximum number of ListPermissions requests that you can make, per second, in this account in the current region. |
| Rate of ListWorkspaces requests | Each supported Region: 5 per second | No | The maximum number of ListWorkspaces requests that you can make, per second, in this account in the current region. |
| Rate of UpdatePermissions requests | Each supported Region: 10 per second | No | The maximum number of UpdatePermissions requests that you can make, per second, in this account in the current region. |
| Rate of UpdateWorkspace requests | Each supported Region: 10 per second | No | The maximum number of UpdateWorkspace requests that you can make, per second, in this account in the current region. |
| Rate of UpdateWorkspaceAuthentication requests | Each supported Region: 1 per second | No | The maximum number of UpdateWorkspaceAuthentication requests that you can make, per second, in this account in the current region. |

Additionally, Amazon Managed Grafana has the following quotas within each workspace

| Resource | Adjustable | Default Quota |
| --- | --- | --- |
| Alerts<br><br>The number of rules per workspace in classic alerting, or the number of rule instances per workspace in Grafana alerting. | No | 100 per workspace. |
| Alert evaluation timeout | No | 30 seconds. |
| Dashboards | No | 2,000 per workspace. |
| Data sources | No | 2,000 per workspace. |
| Data source timeout | No | 60 seconds. |
| Users | No | 10,000 provisioned, 500 concurrent per workspace. |
| API keys | No | 100 per workspace. |
| Service accounts | No | 100 per workspace. |
| Service account tokens<br><br>Active tokens and expired tokens count toward this quota. Delete tokens to remove them from the quota. | No | 100 per workspace. |
| Network access control: Prefix lists | No | 5 per workspace. |
| Network access control: IP address ranges | No | 100 per prefix list. |

| Resource | Adjustable | Default Quota |
| --- | --- | --- |
| Network access control: VPC endpoints | No | 5 per workspace. |

# AWS Ground Station endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | groundstation.us-east-2.amazonaws.com | HTTPS |
| | | groundstation-fips.us-east-2.api.aws | HTTPS |
| | | groundstation-fips.us-east-2.amazonaws.com | HTTPS |
| | | groundstation.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | groundstation.us-east-1.amazonaws.com | HTTPS |
| | | groundstation-fips.us-east-1.api.aws | HTTPS |
| | | groundstation-fips.us-east-1.amazonaws.com | HTTPS |
| | | groundstation.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | groundstation.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | groundstation-fips.us-west-2.api.aws | HTTPS | |
| | | groundstation-fips.us-west-2.amazonaws.com | HTTPS | |
| | | groundstation.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | groundstation.af-south-1.amazonaws.com | HTTPS | |
| | | groundstation.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | groundstation.ap-northeast-2.amazonaws.com | HTTPS | |
| | | groundstation.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | groundstation.ap-southeast-1.amazonaws.com | HTTPS | |
| | | groundstation.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | groundstation.ap-southeast-2.amazonaws.com | HTTPS | |
| | | groundstation.ap-southeast-2.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | groundstation.eu-central-1.amazonaws.com | HTTPS | |
| | | groundstation.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | groundstation.eu-west-1.amazonaws.com | HTTPS | |
| | | groundstation.eu-west-1.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | groundstation.eu-north-1.amazonaws.com | HTTPS | |
| | | groundstation.eu-north-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | groundstation.me-south-1.amazonaws.com<br>groundstation.me-south-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | groundstation.sa-east-1.amazonaws.com<br>groundstation.sa-east-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Config limit | Each supported Region: 100 | Yes | The maximum number of configs allowed. |
| Contact Lead Time Maximum | Each supported Region: 7 | Yes | Maximum lead time allowed for scheduling a contact in days |
| Dataflow endpoint group limit | Each supported Region: 100 | Yes | The maximum number of dataflow endpoint groups allowed. |
| Dataflow endpoints per group limit | Each supported Region: 20 | Yes | The maximum number of dataflow endpoints per group allowed. |
| Enabled Ephemerides limit | Each supported Region: 30 | Yes | Maximum number of enabled customer-provided ephemerides per satellite. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Ephemeris Validation limit | Each supported Region: 10 | Yes | Maximum number of customer-provided ephemerides that be validated simultaneously. |
| Maximum Contact Duration | Each supported Region: 20 | Yes | The maximum contact duration permitted in minutes |
| Mission profile limit | Each supported Region: 100 | Yes | The maximum number of mission profiles allowed. |
| Scheduled Contacts Limit | Each supported Region: 100 | Yes | Maximum number of scheduled contacts allowed |
| Scheduled Minutes Limit | Each supported Region: 1,000 | Yes | The maximum number of scheduled minutes allowed |

## Amazon GuardDuty endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | guardduty.us-east-2.amazonaws.com | HTTPS | |
| | | guardduty-fips.us-east-2.amazonaws.com | HTTPS | |
| | | guardduty.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | guardduty.us-east-1.amazonaws.com | HTTPS | |
| | | guardduty-fips.us-east-1.amazonaws.com | HTTPS | |
| | | guardduty.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | guardduty.us-west-1.amazonaws.com | HTTPS | |
| | | guardduty-fips.us-west-1.amazonaws.com | HTTPS | |
| | | guardduty.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | guardduty.us-west-2.amazonaws.com | HTTPS | |
| | | guardduty-fips.us-west-2.amazonaws.com | HTTPS | |
| | | guardduty.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | guardduty.af-south-1.amazonaws.com | HTTPS | |
| | | guardduty.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | guardduty.ap-east-1.amazonaws.com | HTTPS | |
| | | guardduty.ap-east-1.api.aws | HTTPS | |
| Asia Pacific | ap-south-2 | guardduty.ap-south-2.amazonaws.com | HTTPS | |
| | | guardduty.ap-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | guardduty.ap-southeast-3.amazonaws.com<br>guardduty.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | guardduty.ap-southeast-5.amazonaws.com<br>guardduty.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | guardduty.ap-southeast-4.amazonaws.com<br>guardduty.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | guardduty.ap-south-1.amazonaws.com<br>guardduty.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | guardduty.ap-southeast-6.amazonaws.com<br>guardduty.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | guardduty.ap-northeast-3.amazonaws.com<br>guardduty.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | guardduty.ap-northeast-2.amazonaws.com<br>guardduty.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | guardduty.ap-southeast-1.amazonaws.com | HTTPS | |
| | | guardduty.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | guardduty.ap-southeast-2.amazonaws.com | HTTPS | |
| | | guardduty.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | guardduty.ap-east-2.amazonaws.com | HTTPS | |
| | | guardduty.ap-east-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | guardduty.ap-southeast-7.amazonaws.com | HTTPS | |
| | | guardduty.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | guardduty.ap-northeast-1.amazonaws.com | HTTPS | |
| | | guardduty.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | guardduty.ca-central-1.amazonaws.com | HTTPS | |
| | | guardduty-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | guardduty.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | guardduty.ca-west-1.amazonaws.com | HTTPS | |
| | | guardduty-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | guardduty.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | guardduty.eu-central-1.amazonaws.com | HTTPS | |
| | | guardduty.eu-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | guardduty.eu-west-1.amazonaws.com | HTTPS | |
| | | guardduty.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | guardduty.eu-west-2.amazonaws.com | HTTPS | |
| | | guardduty.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | guardduty.eu-south-1.amazonaws.com | HTTPS | |
| | | guardduty.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | guardduty.eu-west-3.amazonaws.com | HTTPS | |
| | | guardduty.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | guardduty.eu-south-2.amazonaws.com | HTTPS | |
| | | guardduty.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | guardduty.eu-north-1.amazonaws.com | HTTPS | |
| | | guardduty.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | guardduty.eu-central-2.amazonaws.com | HTTPS | |
| | | guardduty.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | guardduty.il-central-1.amazonaws.com | HTTPS | |
| | | guardduty.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | guardduty.mx-central-1.amazonaws.com | HTTPS | |
| | | guardduty.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | guardduty.me-south-1.amazonaws.com<br>guardduty.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | guardduty.me-central-1.amazonaws.com<br>guardduty.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | guardduty.sa-east-1.amazonaws.com<br>guardduty.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | guardduty.us-gov-east-1.amazonaws.com<br>guardduty.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | guardduty.us-gov-west-1.amazonaws.com<br>guardduty.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Detectors | Each supported Region: 1 | No | The maximum number of detector resources that you can create per AWS account per region. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Filters | Each supported Region: 100 | No | The maximum number of saved filters per AWS account per region. |
| Finding retention period | Each supported Region: 90 | No | The maximum number of days a finding is retained. After 90 days findings are deleted. |
| Member accounts by invitation | Each supported Region: 5,000 | No | The maximum number of member accounts that can be associated with a GuardDuty administrator account by invitation. |
| Member accounts through AWS Organizations | Each supported Region: 50,000 | No | Your current AWS Organizations member account quota shows the default maximum number of member accounts that you can associate with an administrator through AWS Organizations, including members added by invitation. The number of GuardDuty member accounts added through AWS Organizations cant exceed the total member accounts in your organization. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Threat intel sets | Each supported Region: 6 | No | The maximum number of Threat intel sets that you can add per AWS account per region. |
| Trusted IP sets | Each supported Region: 1 | No | The maximum number of Trusted IP sets that you can add per AWS account per region. |

# AWS Health endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | health.us-east-2.amazonaws.com | HTTPS |
| | | global.health.amazonaws.com | HTTPS |
| | | health-fips.us-east-2.amazonaws.com | HTTPS |
| | | health.us-east-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | health.us-east-1.amazonaws.com | HTTPS |
| | | global.health.amazonaws.com | HTTPS |
| | | health-fips.us-east-1.amazonaws.com | HTTPS |
| | | health.us-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | health.us-gov-east-1.amazonaws.com | HTTPS |
| | | health-fips.us-gov-east-1.api.aws | HTTPS |
| | | health-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | health.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | health.us-gov-west-1.amazonaws.com | HTTPS |
| | | health-fips.us-gov-west-1.api.aws | HTTPS |
| | | health-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | health.us-gov-west-1.api.aws | HTTPS |

For more information, see Accessing the AWS Health API in the *AWS Health User Guide*.

### Service quotas

This service has no quotas.

## AWS HealthImaging endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | medical-imaging.us-east-1.amazonaws.com | HTTPS |
| | | medical-imaging-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | medical-imaging.us-west-2.amazonaws.com | HTTPS |
| | | medical-imaging-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | medical-imaging.ap-southeast-2.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | medical-imaging.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | medical-imaging.eu-west-2.amazonaws.com | HTTPS |

If you are using HTTP requests to call AWS HealthImaging API actions, you must use two different endpoints depending on the actions being called. The following menu lists the available service endpoints for HTTP requests and the actions they support.

**Supported API actions for HTTP requests**

Using HTTP requests, the following *data store*, *import*, and *tagging* actions are accessible via endpoint:

`https://medical-imaging.region.amazonaws.com`

- CreateDatastore
- GetDatastore
- ListDatastores
- DeleteDatastore
- StartDICOMImportJob
- GetDICOMImportJob
- ListDICOMImportJobs
- TagResource
- ListTagsForResource
- UntagResource

Using HTTP requests, the following *runtime* actions are accessible via endpoint:

`https://runtime-medical-imaging.region.amazonaws.com`

- SearchImageSets
- GetImageSet
- GetImageSetMetadata
- GetImageFrame
- ListImageSetVersions

- UpdateImageSetMetadata
- CopyImageSet
- DeleteImageSet

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Maximum concurrent CopyImageSet requests per data store | Each supported Region: 100 | Yes | The maximum concurrent CopyImageSet requests per data store in the current AWS Region |
| Maximum concurrent DeleteImageSet requests per data store | Each supported Region: 100 | Yes | The maximum concurrent DeleteImageSet requests per data store in the current AWS Region |
| Maximum concurrent UpdateImageSetMetadata requests per data store | Each supported Region: 100 | Yes | The maximum concurrent UpdateImageSetMetadata requests per data store in the current AWS Region |
| Maximum concurrent import jobs per data store | ap-southeast-2: 20<br><br>Each of the other supported Regions: 100 | Yes | The maximum number of concurrent import jobs per data store in the current AWS Region |
| Maximum data stores | Each supported Region: 100 | Yes | The maximum number of active data stores in the current AWS Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of ImageFrames allowed to be copied per CopyImageSet request | Each supported Region: 1,000 | Yes | The maximum number of ImageFrames allowed to be copied per CopyImageSet request in the current AWS Region |
| Maximum number of files in a DICOM import job | Each supported Region: 5,000 | Yes | The maximum number of files in a DICOM import job in the current AWS Region |
| Maximum number of nested folders in a DICOM import job | Each supported Region: 10,000 | No | The maximum number of nested folders in a DICOM import job in the current AWS Region |
| Maximum payload size limit (in KB) accepted by UpdateImageSetMetadata | Each supported Region: 10 Kilobytes | Yes | The maximum payload size limit (in KB) accepted by UpdateImageSetMetadata in the current AWS Region |
| Maximum size (in GB) of all files in a DICOM import job | Each supported Region: 10 Gigabytes | No | The maximum size (in GB) of all files in a DICOM import job in the current AWS Region |
| Maximum size (in GB) of each DICOM P10 file in a DICOM import job | Each supported Region: 10 Gigabytes | No | The maximum size (in GB) of each DICOM P10 file in the DICOM import job in the current AWS Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum size limit (in MB) on ImageSetMetadata per Import, Copy, and UpdateImageSet | Each supported Region: 50 Megabytes | Yes | The maximum size limit (in MB) on ImageSetMetadata per Import, Copy, and UpdateImageSet in the current AWS Region |

# AWS HealthLake endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | healthlake.us-east-2.amazonaws.com | HTTPS |
| | | healthlake-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | healthlake.us-east-1.amazonaws.com | HTTPS |
| | | healthlake-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | healthlake.us-west-2.amazonaws.com | HTTPS |
| | | healthlake-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | healthlake.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | healthlake.ap-southeast-2.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | healthlake.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | healthlake.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | healthlake.eu-west-2.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| C-CDA files per transformation job | Each supported Region: 1,000 | Yes | The maximum number of C-CDA source files in a single bulk transformation job. |
| CSV files per transformation job | Each supported Region: 20 | No | The maximum number of CSV source files in a single bulk transformation job. |
| Columns per CSV file in transformation jobs | Each supported Region: 100 | No | The maximum number of columns in a CSV source |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | file for transformation jobs. |
| Concurrent transformation jobs per account | Each supported Region: 1 | Yes | The maximum number of transformation jobs that can run concurrently per account. |
| Maximum C-CDA file size for sync conversion | Each supported Region: 1 Megabytes | Yes | The maximum size of a C-CDA input in a synchronous conversion request. |
| Maximum CSV files per sync conversion | Each supported Region: 20 | No | The maximum number of CSV files in a single synchronous (real-time) conversion request. |
| Maximum combined CSV file size for sync conversion | Each supported Region: 500 Kilobytes | Yes | The maximum combined size of all CSV files in a single synchronous conversion request. |
| Number of active Subscription per account | Each supported Region: 100 | Yes | The maximum number of active Subscription resources per account. |
| Number of active Subscription resources per datastore | Each supported Region: 50 | Yes | The maximum number of active Subscription resources per datastore. |
| Number of active SubscriptionTopic per account | Each supported Region: 100 | Yes | The maximum number of active SubscriptionTopic resources per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of active SubscriptionTopic resources per datastore | Each supported Region: 50 | Yes | The maximum number of active SubscriptionTopic resources per datastore. |
| Number of characters in a medical note | Each supported Region: 10,000 | No | The maximum number of characters in an individual medical note within the DocumentReference Resource type (POST/PUT requests). |
| Number of concurrent StartFHIR ExportJob jobs | Each supported Region: 1 | No | The maximum concurrent StartFHIRExportJob jobs. |
| Number of concurrent StartFHIR ImportJob jobs | Each supported Region: 1 | No | The maximum concurrent StartFHIRImportJob jobs. |
| Number of data stores per account | Each supported Region: 20 | Yes | The default maximum number of active data stores per account. |
| Number of files in a StartFHIR ImportJob | Each supported Region: 1,000,000 | Yes | The maximum number of files in a StartFHIR ImportJob. |
| Number of published versions per profile | Each supported Region: 99 | No | The maximum number of versions that a profile can have. |
| Number of resources per Bundle | Each supported Region: 500 | No | The maximum number of resources allowed in a Bundle request. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of transformation profiles per account | Each supported Region: 20 | No | The maximum number of transformation profiles per account. |
| Rate of CancelFHIRExportJob requests using DELETE per account | Each supported Region: 1 | No | The maximum number of CancelFHIRExportJob requests using DELETE that you can make per second per account. |
| Rate of CreateDataTransformationProfile requests per account | Each supported Region: 1 | Yes | The maximum number of CreateDataTransformationProfile requests that you can make per second per account. |
| Rate of CreateFHIRDatastore requests per account | Each supported Region: 1 | No | The maximum number of CreateFHIRDatastore requests that you can make per minute per account. |
| Rate of DeleteDataTransformationProfile requests per account | Each supported Region: 1 | Yes | The maximum number of DeleteDataTransformationProfile requests that you can make per second per account. |
| Rate of DeleteFHIRDatastore requests per account | Each supported Region: 1 | No | The maximum number of DeleteFHIRDatastore requests that you can make per minute per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DescribeDataTransformationJob requests per account | Each supported Region: 1 | Yes | The maximum number of DescribeDataTransformationJob requests that you can make per second per account. |
| Rate of DescribeFHIRBulkDeleteJob requests per account | Each supported Region: 10 | Yes | The maximum number of DescribeFHIRBulkDeleteJob requests that you can make per second per account. |
| Rate of DescribeFHIRBulkDeleteJob requests per data store | Each supported Region: 10 | Yes | The maximum number of DescribeFHIRBulkDeleteJob requests that you can make per second per data store. |
| Rate of DescribeFHIRDatastore requests per account | Each supported Region: 10 | No | The maximum number of DescribeFHIRDatastore requests that you can make per second per account. |
| Rate of DescribeFHIRExportJob requests per account | Each supported Region: 10 | No | The maximum number of DescribeFHIRExportJob requests that you can make per second per account. |
| Rate of DescribeFHIRExportJob requests using GET per account | Each supported Region: 10 | No | The maximum number of DescribeFHIRExportJob requests using GET that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeFHIRImportJob requests per account | Each supported Region: 10 | No | The maximum number of DescribeFHIRImportJob requests that you can make per second per account. |
| Rate of Discovery requests per account | Each supported Region: 10 | No | The maximum number of Discovery requests that you can make per minute per account. |
| Rate of GET requests per account | Each supported Region: 6,000 | Yes | The maximum number of GET requests that you can make per second per account. |
| Rate of GET requests per data store | Each supported Region: 3,000 | Yes | The maximum number of GET requests that you can make per second per data store. Data stores created prior to 8/21/2023 will be limited to 100 requests per second. |
| Rate of GetCapabilities requests per account | Each supported Region: 10 | No | The maximum number of GetCapabilities requests that you can make per second per account. |
| Rate of GetDataTransformationProfile requests per account | Each supported Region: 1 | Yes | The maximum number of GetDataTransformationProfile requests that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetExportedFile requests per datastore | Each supported Region: 10 | No | The maximum number of GetExportedFile requests that you can make per second per datastore. |
| Rate of ListDataTransformationJobs requests per account | Each supported Region: 1 | Yes | The maximum number of ListDataTransformationJobs requests that you can make per second per account. |
| Rate of ListDataTransformationProfileVersions requests per account | Each supported Region: 1 | Yes | The maximum number of ListDataTransformationProfileVersions requests that you can make per second per account. |
| Rate of ListDataTransformationProfiles requests per account | Each supported Region: 1 | Yes | The maximum number of ListDataTransformationProfiles requests that you can make per second per account. |
| Rate of ListFHIRDatastores requests per account | Each supported Region: 10 | No | The maximum number of ListFHIRDatastores requests that you can make per second per account. |
| Rate of ListFHIRExportJobs requests per account | Each supported Region: 10 | No | The maximum number of ListFHIRExportJobs requests that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListFHIRImportJobs requests per account | Each supported Region: 10 | No | The maximum number of ListFHIRImportJobs requests that you can make per second per account. |
| Rate of ListTagsforResource requests per account | Each supported Region: 10 | No | The maximum number of ListTagsforResource requests that you can make per second per account. |
| Rate of PublishDataTransformationProfile requests per account | Each supported Region: 1 | Yes | The maximum number of PublishDataTransformationProfile requests that you can make per second per account. |
| Rate of SearchEverything requests per account | Each supported Region: 10 | Yes | The maximum number of SearchEverything requests that you can make per second per account. |
| Rate of SearchEverything requests per data store | Each supported Region: 10 | Yes | The maximum number of SearchEverything requests that you can make per second per data store. |
| Rate of StartDataTransformationJob requests per account | Each supported Region: 1 | Yes | The maximum number of StartDataTransformationJob requests that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartFHIRBulkDeleteJob requests per account | Each supported Region: 10 | Yes | The maximum number of StartFHIRBulkDeleteJob requests that you can make per second per account. |
| Rate of StartFHIRBulkDeleteJob requests per data store | Each supported Region: 10 | Yes | The maximum number of StartFHIRBulkDeleteJob requests that you can make per second per data store. |
| Rate of StartFHIRBulkMemberMatchJob requests per account | Each supported Region: 10 | Yes | The maximum number of StartFHIRBulkMemberMatchJob requests that you can make per second per account. |
| Rate of StartFHIRBulkMemberMatchJob requests per data store | Each supported Region: 10 | Yes | The maximum number of StartFHIRBulkMemberMatchJob requests that you can make per second per data store. |
| Rate of StartFHIRExportJob requests per account | Each supported Region: 1 | No | The maximum number of StartFHIRExportJob requests that you can make per second per account. |
| Rate of StartFHIRExportJob requests using GET per account | Each supported Region: 1 | No | The maximum number of StartFHIRExportJob requests using GET that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartFHIRExportJob requests using POST per account | Each supported Region: 1 | No | The maximum number of StartFHIRExportJob requests using POST that you can make per second per account. |
| Rate of StartFHIRImportJob requests per account | Each supported Region: 25 | No | The maximum number of StartFHIRImportJob requests that you can make per second per account. |
| Rate of TagResource requests per account | Each supported Region: 10 | No | The maximum number of TagResource requests that you can make per second per account. |
| Rate of UntagResource requests per account | Each supported Region: 10 | No | The maximum number of UntagResource requests that you can make per second per account. |
| Rate of UpdateDataTransformationProfile requests per account | Each supported Region: 1 | Yes | The maximum number of UpdateDataTransformationProfile requests that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ValidateResource requests per account | Each supported Region: 2,000 | Yes | The maximum number of ValidateResource requests that you can make per second per account. Data stores created prior to 8/21/2023 will be limited to 300 requests per second. |
| Rate of ValidateResource requests per data store | Each supported Region: 1,000 | Yes | The maximum number of ValidateResource requests that you can make per second per data store. Data stores created prior to 8/21/2023 will be limited to 300 requests per second. |
| Rate of ValidateSource per account | Each supported Region: 1 | Yes | The maximum number of ValidateSource requests that you can make per second per account. |
| Rate of WRITE requests per account | Each supported Region: 6,000 | Yes | The maximum number of CREATE\|UPDATE\|DELETE requests that you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of WRITE requests per data store | Each supported Region: 3,000 | Yes | The maximum number of CREATE\|UPDATE\|DELETE requests that you can make per second per data store. Data stores created prior to 8/21/2023 will be limited to 300 requests per second. |
| Rate of search requests using GET per account | Each supported Region: 200 | Yes | The maximum number of search requests using GET that you can make per second per account. |
| Rate of search requests using GET per data store | Each supported Region: 100 | Yes | The maximum number of search requests using GET that you can make per second per data store. |
| Rate of search requests using POST per account | Each supported Region: 200 | Yes | The maximum number of search requests using POST that you can make per second per account. |
| Rate of search requests using POST per data store | Each supported Region: 100 | Yes | The maximum number of search requests using POST that you can make per second per data store. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of sync (real-time) conversion requests (TransformData) per account | Each supported Region: 1 | Yes | The maximum number of TransformData requests that you can make per second per account. |
| Size of individual imported file | Each supported Region: 50 Gigabytes | No | The maximum size (in GB) of an individual file included in a StartFHIR ImportJob. |
| Size per CSV file in a transformation job | Each supported Region: 50 Megabytes | Yes | The maximum size of an individual CSV source file in a bulk transformation job. |
| Total Queued Async Bundle Transactions per datastore | Each supported Region: 500 | Yes | The maximum number of queued async bundle transactions per datastore at any given time. |
| Total Queued Bulk Export jobs per datastore | Each supported Region: 25 | Yes | The maximum number of queued bulk export jobs per datastore at any given time. |
| Total Queued Bulk Import jobs per datastore | Each supported Region: 25 | Yes | The maximum number of queued bulk import jobs per datastore at any given time. |
| Total import job size | Each supported Region: 5,000 Gigabytes | Yes | The maximum size (in GB) of all files included in the import job. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Total input size per transformation job | Each supported Region: 1 Gigabytes | Yes | The maximum combined size of all source files in a single bulk transformation job. |

# Service quotas and endpoints for AWS HealthOmics

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

HealthOmics provides separate API endpoints for the storage, analytics, and workflow services. The API quotas section in the *HealthOmics User Guide* provides the list of API operations for each service. The API endpoints use the following service codes:

| Service | Service code |
| --- | --- |
| HealthOmics Workflows | workflows |
| HealthOmics Storage | storage |
| HealthOmics Analytics | analytics |

The resulting API endpoints have the following structure:

```
workflows-omics.region.amazonaws.com
storage-omics.region.amazonaws.com
```

```
analytics-omics.region.amazonaws.com
```

For regions that support FIPS API endpoints, the FIPS endpoints have the following structure:

```
workflows-omics-fips.region.amazonaws.com
storage-omics-fips.region.amazonaws.com
analytics-omics-fips.region.amazonaws.com
```

The following table lists the API endpoints for Regions where the HealthOmics service is available.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | **Workflows:** workflows-omics.us-east-1.amazonaws.com<br><br>**Storage:** storage-omics.us-east-1.amazonaws.com<br><br>**Analytics:** analytics-omics.us-east-1.amazonaws.com<br><br>**FIPS**<br><br>**Workflows:** workflows-omics-fips.us-east-1.amazonaws.com<br><br>**Storage:** storage-omics-fips.us-east-1.amazonaws.com<br><br>**Analytics:** analytics-omics-fips.us-east-1.amazonaws.com | HTTPS |
| US East (Ohio) | us-east-2 | **Workflows:** workflows-omics.us-east-2.amazonaws.com<br><br>**FIPS**<br><br>**Workflows:** workflows-omics-fips.us-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | **Workflows:** workflows-omics.us-west-2.amazonaws.com<br><br>**Storage:** storage-omics.us-west-2.amazonaws.com<br><br>**Analytics:** analytics-omics.us-west-2.amazonaws.com<br><br>**FIPS**<br><br>**Workflows:** workflows-omics-fips.us-west-2.amazonaws.com<br><br>**Storage:** storage-omics-fips.us-west-2.amazonaws.com<br><br>**Analytics:** analytics-omics-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | **Workflows:** workflows-omics.ap-southeast-1.amazonaws.com<br><br>**Storage:** storage-omics.ap-southeast-1.amazonaws.com<br><br>**Analytics:** analytics-omics.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | **Workflows:** workflows-omics.ap-northeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | **Workflows:** workflows-omics.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | **Workflows:** workflows-omics.eu-central-1.amazonaws.com<br><br>**Storage:** storage-omics.eu-central-1.amazonaws.com<br><br>**Analytics:** analytics-omics.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | **Workflows:** workflows-omics.eu-west-1.amazonaws.com<br><br>**Storage:** storage-omics.eu-west-1.amazonaws.com<br><br>**Analytics:** analytics-omics.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | **Workflows:** workflows-omics.eu-west-2.amazonaws.com<br><br>**Storage:** storage-omics.eu-west-2.amazonaws.com<br><br>**Analytics:** analytics-omics.eu-west-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | **Workflows:** workflows-omics.il-central-1.amazonaws.com<br><br>**Storage:** storage-omics.il-central-1.amazonaws.com<br><br>**Analytics:** analytics-omics.il-central-1.amazonaws.com | HTTPS |

# Quotas for AWS HealthOmics

Your AWS account initially contains the HealthOmics service default quotas in each Region. You can request increases for some quotas, and other quotas cannot be increased.

To view the quotas, open the Service Quotas console. In the navigation pane, choose **AWS services** and select **HealthOmics**.

To request a quota increase, see Requesting a Quota Increase in the *Service Quotas User Guide*. If the quota is not yet available in Service Quotas, use the limit increase form.

The table below shows the default values for each of the HealthOmics quotas.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Analytics - Maximum annotation stores | Each supported Region: 10 | Yes | The maximum number of annotation stores in the current AWS region |
| Analytics - Maximum concurrent variant or annotation store import jobs | Each supported Region: 5 | Yes | The maximum number of concurrent import jobs in the current AWS region |
| Analytics - Maximum files per variant store import job | Each supported Region: 1,000 | Yes | The maximum number of files per variant import job in the current AWS region |
| Analytics - Maximum shares per annotation store | Each supported Region: 10 | Yes | The maximum number of shares per annotation store in the current AWS region |
| Analytics - Maximum shares per variant store | Each supported Region: 10 | Yes | The maximum number of shares per variant store in the current AWS region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Analytics - Maximum size of each file in a variant import job | Each supported Region: 20 Gigabytes | Yes | The maximum size of one file in a variant import job in the current AWS region |
| Analytics - Maximum size of each file in an annotation import job | Each supported Region: 20 Gigabytes | Yes | The maximum size of one file in an annotation import job in the current AWS region |
| Analytics - Maximum variant stores | Each supported Region: 10 | Yes | The maximum number of variant stores in the current AWS region |
| Analytics - Maximum versions per annotation store | Each supported Region: 10 | Yes | The maximum number of versions per annotation store in the current AWS region |
| Configurations - Maximum configurations | Each supported Region: 10 | Yes | The maximum number of configurations in the current AWS region. |
| Storage - Maximum concurrent read set activation jobs | Each supported Region: 25 | Yes | The maximum number of concurrent read set activation jobs in the current AWS region |
| Storage - Maximum concurrent sequence and reference store export jobs | Each supported Region: 5 | Yes | The maximum number of concurrent export jobs from a sequence or reference store in the current AWS region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Storage - Maximum concurrent sequence or reference store import jobs | Each supported Region: 5 | Yes | The maximum number of concurrent import jobs for a sequence or reference store in the current AWS region |
| Storage - Maximum read sets per sequence store | Each supported Region: 1,000,000 | Yes | The maximum number of read sets in a sequence store in the current AWS region |
| Storage - Maximum references per reference store | Each supported Region: 50 | Yes | The maximum number of references in a reference store in the current AWS region |
| Storage - Maximum sequence stores | Each supported Region: 20 | Yes | The maximum number of sequence stores in the current AWS region |
| Workflows - Maximum active GPUs | Each supported Region: 12 | Yes | The maximum number of concurrent active GPUs in the current AWS region. In us-east-1 and us-west-2, quota increase requests for values up to 500 are automatically approved. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Workflows - Maximum concurrent active runs using dynamic run storage | Each supported Region: 50 | Yes | The maximum number of active runs using dynamic run storage in the current AWS region. Quota increase requests for values up to 200 are automatically approved. |
| Workflows - Maximum concurrent active runs using static run storage | Each supported Region: 10 | Yes | The maximum number of active runs using static run storage in the current AWS region. Quota increase requests for values up to 50 are automatically approved. |
| Workflows - Maximum concurrent tasks per run | Each supported Region: 25 | Yes | The maximum number of concurrent tasks in each run in the current AWS region. In us-east-1 and us-west-2, quota increase requests for values up to 100 are automatically approved. |
| Workflows - Maximum run duration | Each supported Region: 604,800 Seconds | Yes | The maximum workflow run duration in the current AWS region. |
| Workflows - Maximum runs (active or inactive) | Each supported Region: 100,000 | Yes | The maximum number of runs (active or inactive) in the current AWS region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Workflows - Maximum shares per workflow | Each supported Region: 100 | Yes | The maximum number of shares per workflow in the current AWS region |
| Workflows - Maximum static run storage capacity per run | Each supported Region: 9,600 | Yes | The maximum static run storage capacity in gibibytes (GiB) for each run in the current AWS region. In us-east-1 and us-west-2, quota increase requests for values up to 50,000 are automatically approved. |
| Workflows - Maximum workflows | Each supported Region: 1,000 | Yes | The maximum number of workflows in the current AWS region. |
| Workflows - Transactions per second (TPS) for the StartRun operation | Each supported Region: 5 | Yes | The maximum transactions per second (TPS) for the StartRun operation in the current AWS region. Quota increase requests for values up to 10 are automatically approved. |

For information about API operation quotas, see Quotas in the *HealthOmics User Guide*.

## AWS Identity and Access Management endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | iam.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | iam.amazonaws.com<br>iam-fips.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | iam.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | iam.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | iam.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | iam.amazonaws.com | HTTPS |
| Asia Pacific | ap-south-2 | iam.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | iam.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | iam.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | iam.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | iam.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | iam.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | iam.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | iam.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | iam.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | iam.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | iam.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | iam.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | iam.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | iam.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | iam.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | iam.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | iam.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | iam.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | iam.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | iam.us-gov.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | iam.us-gov.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Access keys per user | Each supported Region: 2 | No | The maximum number of access keys that you can create for an IAM user. |
| Customer managed policies per account | Each supported Region: 1,500 | Yes | The maximum number of customer managed policies that you can create in this account. |
| Groups per account | Each supported Region: 300 | Yes | The maximum number of IAM groups that you can create in this account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| IAM groups per user | Each supported Region: 10 | No | The maximum number of IAM groups to which you can add an IAM user. |
| Identity providers per IAM SAML provider object | Each supported Region: 10 | No | The maximum number of identity providers (IdPs) that you can add to an IAM SAML provider object. |
| Instance profiles per account | Each supported Region: 1,000 | Yes | The maximum number of instance profiles that you can create in this account. |
| Keys per SAML provider | Each supported Region: 10 | No | The maximum number of keys that you can assign to a SAML provider. |
| MFA devices per user | Each supported Region: 8 | No | The maximum number of MFA devices that you can configure for an IAM user. |
| Managed policies per group | Each supported Region: 10 | No | The maximum number of IAM managed policies that you can attach to an IAM group. |
| Managed policies per role | Each supported Region: 20 | Yes | The maximum number of IAM managed policies that you can attach to an IAM role. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Managed policies per user | Each supported Region: 10 | Yes | The maximum number of IAM managed policies that you can attach to an IAM user. |
| Managed policy length | Each supported Region: 6,144 | No | The maximum number of characters in an IAM managed policy. |
| OpenId connect providers per account | Each supported Region: 100 | Yes | Maximum number of OpenID connectors allowed for an AWS account. |
| Role trust policy length | Each supported Region: 2,048 | Yes | The maximum number of characters in an IAM role trust policy. |
| Roles per account | Each supported Region: 1,000 | Yes | The maximum number of IAM roles that you can create in this account. |
| SAML providers per account | Each supported Region: 100 | No | The maximum number of SAML providers that you can create in this account. |
| SSH Public keys per user | Each supported Region: 5 | No | The maximum number of SSH public keys that you can assign to an IAM user. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Server certificates per account | Each supported Region: 20 | Yes | The maximum number of server certificates that you can store in this account. |
| Signing certificates per user | Each supported Region: 2 | No | The maximum number of signing certificates that you can upload for an IAM user. |
| Tags per role | Each supported Region: 50 | No | The maximum number of tags that you can assign to an IAM role. |
| Tags per user | Each supported Region: 50 | No | The maximum number of tags that you can assign to an IAM user. |
| Users per account | Each supported Region: 5,000 | No | The maximum number of IAM users you can create for your AWS account. |
| Versions per managed policy | Each supported Region: 5 | No | The maximum number of versions that you can save to an IAM managed policy in this account before you must overwrite an existing version. |

For more information about IAM quotas, see IAM and AWS STS quotas in the *IAM User Guide*.

# IAM Access Analyzer endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | access-analyzer.us-east-2.amazonaws.com | HTTPS |
| | | access-analyzer-fips.us-east-2.api.aws | HTTPS |
| | | access-analyzer-fips.us-east-2.amazonaws.com | HTTPS |
| | | access-analyzer.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | access-analyzer.us-east-1.amazonaws.com | HTTPS |
| | | access-analyzer-fips.us-east-1.api.aws | HTTPS |
| | | access-analyzer-fips.us-east-1.amazonaws.com | HTTPS |
| | | access-analyzer.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | access-analyzer.us-west-1.amazonaws.com | HTTPS |
| | | access-analyzer-fips.us-west-1.api.aws | HTTPS |
| | | access-analyzer-fips.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | access-analyzer.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | access-analyzer.us-west-2.amazonaws.com<br>access-analyzer-fips.us-west-2.api.aws<br>access-analyzer-fips.us-west-2.amazonaws.com<br>access-analyzer.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | access-analyzer.af-south-1.amazonaws.com<br>access-analyzer.af-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | access-analyzer.ap-east-1.amazonaws.com<br>access-analyzer.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | access-analyzer.ap-south-2.amazonaws.com<br>access-analyzer.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | access-analyzer.ap-southeast-3.amazonaws.com<br>access-analyzer.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | access-analyzer.ap-southeast-5.amazonaws.com<br>access-analyzer.ap-southeast-5.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | access-analyzer.ap-southeast-4.amazonaws.com<br>access-analyzer.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | access-analyzer.ap-south-1.amazonaws.com<br>access-analyzer.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | access-analyzer.ap-southeast-6.amazonaws.com<br>access-analyzer.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | access-analyzer.ap-northeast-3.amazonaws.com<br>access-analyzer.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | access-analyzer.ap-northeast-2.amazonaws.com<br>access-analyzer.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | access-analyzer.ap-southeast-1.amazonaws.com<br>access-analyzer.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | access-analyzer.ap-southeast-2.amazonaws.com<br>access-analyzer.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | access-analyzer.ap-east-2.amazonaws.com<br>access-analyzer.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | access-analyzer.ap-southeast-7.amazonaws.com<br>access-analyzer.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | access-analyzer.ap-northeast-1.amazonaws.com<br>access-analyzer.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | access-analyzer.ca-central-1.amazonaws.com<br>access-analyzer-fips.ca-central-1.api.aws<br>access-analyzer-fips.ca-central-1.amazonaws.com<br>access-analyzer.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | access-analyzer.ca-west-1.amazonaws.com<br>access-analyzer-fips.ca-west-1.api.aws<br>access-analyzer-fips.ca-west-1.amazonaws.com<br>access-analyzer.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | access-analyzer.eu-central-1.amazonaws.com<br>access-analyzer.eu-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | access-analyzer.eu-west-1.amazonaws.com | HTTPS |
| | | access-analyzer.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | access-analyzer.eu-west-2.amazonaws.com | HTTPS |
| | | access-analyzer.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | access-analyzer.eu-south-1.amazonaws.com | HTTPS |
| | | access-analyzer.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | access-analyzer.eu-west-3.amazonaws.com | HTTPS |
| | | access-analyzer.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | access-analyzer.eu-south-2.amazonaws.com | HTTPS |
| | | access-analyzer.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | access-analyzer.eu-north-1.amazonaws.com | HTTPS |
| | | access-analyzer.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | access-analyzer.eu-central-2.amazonaws.com | HTTPS |
| | | access-analyzer.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | access-analyzer.il-central-1.amazonaws.com | HTTPS |
| | | access-analyzer.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | access-analyzer.mx-central-1.amazonaws.com | HTTPS |
| | | access-analyzer.mx-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | access-analyzer.me-south-1.amazonaws.com<br>access-analyzer.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | access-analyzer.me-central-1.amazonaws.com<br>access-analyzer.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | access-analyzer.sa-east-1.amazonaws.com<br>access-analyzer.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | access-analyzer.us-gov-east-1.amazonaws.com<br>access-analyzer.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | access-analyzer.us-gov-west-1.amazonaws.com<br>access-analyzer.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Access previews per analyzer per hour | Each supported Region: 1,000 | Yes | The maximum number of access previews per analyzer per hour. |
| Account level analyzer | Each supported Region: 1 | No | The maximum number of account level analyzers that can be created per |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | analyzer type per AWS Account per Region. |
| Archive rules per analyzer | Each supported Region: 100 | Yes | The maximum number of archive rules per analyzer. |
| CloudTrail log files processed per policy generation | Each supported Region: 100,000 | No | The maximum number of CloudTrail log files that can be processed per policy generation. |
| Concurrent policy generations | Each supported Region: 1 | No | The maximum number of concurrent policy generations. |
| Organization level analyzer | Each supported Region: 5 | Yes | The maximum number of organization level analyzers that can be created per analyzer type per AWS Account per Region. |
| Policy generation CloudTrail data size | Each supported Region: 25 Gigabytes | No | The maximum size of CloudTrail data per policy generation. |
| Policy generation CloudTrail time range | Each supported Region: 90 | No | The maximum CloudTrail time range that you can select in days when you generate a policy. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Policy generations per day | af-south-1: 5<br>ap-east-1: 5<br>ap-south-2: 5<br>ap-southeast-3: 5<br>ap-southeast-4: 5<br>ca-west-1: 5<br>eu-central-2: 5<br>eu-south-1: 5<br>eu-south-2: 5<br>Each of the other supported Regions: 50 | No | The maximum number of policy generations per day. |

# AWS IAM Identity Center endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## IAM Identity Center

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | sso.us-east-2.amazonaws.com | HTTPS | |
| | | sso-fips.us-east-2.api.aws | HTTPS | |
| | | sso.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | sso.us-east-1.amazonaws.com | HTTPS | |
| | | sso-fips.us-east-1.api.aws | HTTPS | |
| | | sso.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | sso.us-west-1.amazonaws.com | HTTPS | |
| | | sso-fips.us-west-1.api.aws | HTTPS | |
| | | sso.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | sso.us-west-2.amazonaws.com | HTTPS | |
| | | sso-fips.us-west-2.api.aws | HTTPS | |
| | | sso.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | sso.af-south-1.amazonaws.com | HTTPS | |
| | | sso.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | sso.ap-east-1.amazonaws.com | HTTPS | |
| | | sso.ap-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | sso.ap-south-2.amazonaws.com<br>sso.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | sso.ap-southeast-3.amazonaws.com<br>sso.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | sso.ap-southeast-5.amazonaws.com<br>sso.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | sso.ap-southeast-4.amazonaws.com<br>sso.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | sso.ap-south-1.amazonaws.com<br>sso.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | sso.ap-southeast-6.amazonaws.com<br>sso.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | sso.ap-northeast-3.amazonaws.com<br>sso.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | sso.ap-northeast-2.amazonaws.com<br>sso.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | sso.ap-southeast-1.amazonaws.com | HTTPS |
| | | sso.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | sso.ap-southeast-2.amazonaws.com | HTTPS |
| | | sso.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | sso.ap-east-2.amazonaws.com | HTTPS |
| | | sso.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | sso.ap-southeast-7.amazonaws.com | HTTPS |
| | | sso.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | sso.ap-northeast-1.amazonaws.com | HTTPS |
| | | sso.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | sso.ca-central-1.amazonaws.com | HTTPS |
| | | sso-fips.ca-central-1.api.aws | HTTPS |
| | | sso.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | sso.ca-west-1.amazonaws.com | HTTPS |
| | | sso-fips.ca-west-1.api.aws | HTTPS |
| | | sso.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | sso.eu-central-1.amazonaws.com | HTTPS |
| | | sso.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | sso.eu-west-1.amazonaws.com | HTTPS |
| | | sso.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | sso.eu-west-2.amazonaws.com | HTTPS |
| | | sso.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | sso.eu-south-1.amazonaws.com | HTTPS |
| | | sso.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | sso.eu-west-3.amazonaws.com | HTTPS |
| | | sso.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | sso.eu-south-2.amazonaws.com | HTTPS |
| | | sso.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | sso.eu-north-1.amazonaws.com | HTTPS |
| | | sso.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | sso.eu-central-2.amazonaws.com | HTTPS |
| | | sso.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | sso.il-central-1.amazonaws.com | HTTPS |
| | | sso.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | sso.mx-central-1.amazonaws.com | HTTPS |
| | | sso.mx-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | sso.me-south-1.amazonaws.com<br>sso.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | sso.me-central-1.amazonaws.com<br>sso.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | sso.sa-east-1.amazonaws.com<br>sso.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | sso.us-gov-east-1.amazonaws.com<br>sso.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | sso.us-gov-west-1.amazonaws.com<br>sso.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

## Identity Store

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | identitystore.us-east-2.amazonaws.com<br>identitystore.us-east-2.api.aws<br>identitystore-fips.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | identitystore.us-east-1.amazonaws.com | HTTPS |
| | | identitystore.us-east-1.api.aws | HTTPS |
| | | identitystore-fips.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | identitystore.us-west-1.amazonaws.com | HTTPS |
| | | identitystore.us-west-1.api.aws | HTTPS |
| | | identitystore-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | identitystore.us-west-2.amazonaws.com | HTTPS |
| | | identitystore.us-west-2.api.aws | HTTPS |
| | | identitystore-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | identitystore.af-south-1.amazonaws.com | HTTPS |
| | | identitystore.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | identitystore.ap-east-1.amazonaws.com | HTTPS |
| | | identitystore.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | identitystore.ap-south-2.amazonaws.com | HTTPS |
| | | identitystore.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | identitystore.ap-southeast-3.amazonaws.com | HTTPS |
| | | identitystore.ap-southeast-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | identitystore.ap-southeast-5.amazonaws.com<br>identitystore.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | identitystore.ap-southeast-4.amazonaws.com<br>identitystore.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | identitystore.ap-south-1.amazonaws.com<br>identitystore.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | identitystore.ap-southeast-6.amazonaws.com<br>identitystore.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | identitystore.ap-northeast-3.amazonaws.com<br>identitystore.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | identitystore.ap-northeast-2.amazonaws.com<br>identitystore.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | identitystore.ap-southeast-1.amazonaws.com<br>identitystore.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | identitystore.ap-southeast-2.amazonaws.com<br>identitystore.ap-southeast-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | identitystore.ap-east-2.amazonaws.com | HTTPS |
| | | identitystore.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | identitystore.ap-southeast-7.amazonaws.com | HTTPS |
| | | identitystore.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | identitystore.ap-northeast-1.amazonaws.com | HTTPS |
| | | identitystore.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | identitystore.ca-central-1.amazonaws.com | HTTPS |
| | | identitystore.ca-central-1.api.aws | HTTPS |
| | | identitystore-fips.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | identitystore.ca-west-1.amazonaws.com | HTTPS |
| | | identitystore.ca-west-1.api.aws | HTTPS |
| | | identitystore-fips.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | identitystore.eu-central-1.amazonaws.com | HTTPS |
| | | identitystore.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | identitystore.eu-west-1.amazonaws.com | HTTPS |
| | | identitystore.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | identitystore.eu-west-2.amazonaws.com | HTTPS |
| | | identitystore.eu-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | identitystore.eu-south-1.amazonaws.com | HTTPS | |
| | | identitystore.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | identitystore.eu-west-3.amazonaws.com | HTTPS | |
| | | identitystore.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | identitystore.eu-south-2.amazonaws.com | HTTPS | |
| | | identitystore.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | identitystore.eu-north-1.amazonaws.com | HTTPS | |
| | | identitystore.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | identitystore.eu-central-2.amazonaws.com | HTTPS | |
| | | identitystore.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | identitystore.il-central-1.amazonaws.com | HTTPS | |
| | | identitystore.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | identitystore.mx-central-1.amazonaws.com | HTTPS | |
| | | identitystore.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | identitystore.me-south-1.amazonaws.com | HTTPS | |
| | | identitystore.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | identitystore.me-central-1.amazonaws.com | HTTPS | |
| | | identitystore.me-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | identitystore.sa-east-1.amazonaws.com<br>identitystore.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | identitystore.us-gov-east-1.amazonaws.com<br>identitystore.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | identitystore.us-gov-west-1.amazonaws.com<br>identitystore.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## IAM Identity Center OpenID Connect (OIDC)

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | oidc.us-east-2.amazonaws.com<br>oidc-fips.us-east-2.api.aws<br>oidc.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | oidc.us-east-1.amazonaws.com<br>oidc-fips.us-east-1.api.aws<br>oidc.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | oidc.us-west-1.amazonaws.com<br>oidc-fips.us-west-1.api.aws<br>oidc.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | oidc.us-west-2.amazonaws.com | HTTPS |
| | | oidc-fips.us-west-2.api.aws | HTTPS |
| | | oidc.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | oidc.af-south-1.amazonaws.com | HTTPS |
| | | oidc.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | oidc.ap-east-1.amazonaws.com | HTTPS |
| | | oidc.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | oidc.ap-south-2.amazonaws.com | HTTPS |
| | | oidc.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | oidc.ap-southeast-3.amazonaws.com | HTTPS |
| | | oidc.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | oidc.ap-southeast-5.amazonaws.com | HTTPS |
| | | oidc.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | oidc.ap-southeast-4.amazonaws.com | HTTPS |
| | | oidc.ap-southeast-4.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | oidc.ap-south-1.amazonaws.com<br>oidc.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | oidc.ap-southeast-6.amazonaws.com<br>oidc.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | oidc.ap-northeast-3.amazonaws.com<br>oidc.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | oidc.ap-northeast-2.amazonaws.com<br>oidc.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | oidc.ap-southeast-1.amazonaws.com<br>oidc.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | oidc.ap-southeast-2.amazonaws.com<br>oidc.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | oidc.ap-east-2.amazonaws.com<br>oidc.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | oidc.ap-southeast-7.amazonaws.com<br>oidc.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | oidc.ap-northeast-1.amazonaws.com | HTTPS |
| | | oidc.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | oidc.ca-central-1.amazonaws.com | HTTPS |
| | | oidc-fips.ca-central-1.api.aws | HTTPS |
| | | oidc.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | oidc.ca-west-1.amazonaws.com | HTTPS |
| | | oidc-fips.ca-west-1.api.aws | HTTPS |
| | | oidc.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | oidc.eu-central-1.amazonaws.com | HTTPS |
| | | oidc.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | oidc.eu-west-1.amazonaws.com | HTTPS |
| | | oidc.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | oidc.eu-west-2.amazonaws.com | HTTPS |
| | | oidc.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | oidc.eu-south-1.amazonaws.com | HTTPS |
| | | oidc.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | oidc.eu-west-3.amazonaws.com | HTTPS |
| | | oidc.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | oidc.eu-south-2.amazonaws.com | HTTPS |
| | | oidc.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | oidc.eu-north-1.amazonaws.com<br>oidc.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | oidc.eu-central-2.amazonaws.com<br>oidc.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | oidc.il-central-1.amazonaws.com<br>oidc.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | oidc.mx-central-1.amazonaws.com<br>oidc.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | oidc.me-south-1.amazonaws.com<br>oidc.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | oidc.me-central-1.amazonaws.com<br>oidc.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | oidc.sa-east-1.amazonaws.com<br>oidc.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | oidc.us-gov-east-1.amazonaws.com<br>oidc.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | oidc.us-gov-west-1.amazonaws.com<br>oidc.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| File size of service provider SAML 2.0 certificates (in PEM format) | Each supported Region: 2 Kilobytes | No | The maximum file size (in KB) of service provider SAML 2.0 certificates (in PEM format). |
| Number of groups supported in IAM Identity Center | Each supported Region: 100,000 | Yes | The maximum number of groups supported in IAM Identity Center. |
| Number of permission sets allowed in IAM Identity Center | Each supported Region: 3,500 | Yes | The default number of permission sets in IAM Identity Center. |
| Number of permission sets allowed per AWS account | Each supported Region: 500 | Yes | The default number of permission sets allowed per AWS account. |
| Number of users supported in IAM Identity Center | Each supported Region: 200,000 | Yes | The default number of users supported in IAM Identity Center. |
| Total number of AWS accounts or applications that can be configured | Each supported Region: 7,000 | Yes | The default number of AWS accounts or applications that can be configured. This quota applies separately to AWS accounts and to applications. You can configure up to 7000 accounts and up to 7000 applications. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Total number of groups that can be assigned to a permission set per AWS account, or to an application | Each supported Region: 100 | No | Total number of groups that can be assigned to a permission set per AWS account, or to an application. |
| Total number of regions for IAM Identity Center | Each supported Region: 6 | Yes | The total number of regions that can be enabled for IAM Identity Center, including the primary region. |

For more information, see AWS IAM Identity Center quotas in the *AWS IAM Identity Center User Guide*.

## AWS Identity and Access Management Roles Anywhere endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | rolesanywhere.us-east-2.amazonaws.com | HTTPS |
| | | rolesanywhere.us-east-2.api.aws | HTTPS |
| | | rolesanywhere-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | rolesanywhere.us-east-1.amazonaws.com | HTTPS |
| | | rolesanywhere-fips.us-east-1.amazonaws.com | HTTPS |
| | | rolesanywhere.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | rolesanywhere.us-west-1.amazonaws.com | HTTPS |
| | | rolesanywhere.us-west-1.api.aws | HTTPS |
| | | rolesanywhere-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | rolesanywhere.us-west-2.amazonaws.com | HTTPS |
| | | rolesanywhere.us-west-2.api.aws | HTTPS |
| | | rolesanywhere-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | rolesanywhere.af-south-1.amazonaws.com | HTTPS |
| | | rolesanywhere.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | rolesanywhere.ap-east-1.amazonaws.com | HTTPS |
| | | rolesanywhere.ap-east-1.api.aws | HTTPS |
| Asia Pacific | ap-south-2 | rolesanywhere.ap-south-2.amazonaws.com | HTTPS |
| | | rolesanywhere.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rolesanywhere.ap-southeast-3.amazonaws.com<br>rolesanywhere.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rolesanywhere.ap-southeast-5.amazonaws.com<br>rolesanywhere.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rolesanywhere.ap-southeast-4.amazonaws.com<br>rolesanywhere.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rolesanywhere.ap-south-1.amazonaws.com<br>rolesanywhere.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | rolesanywhere.ap-southeast-6.amazonaws.com<br>rolesanywhere.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | rolesanywhere.ap-northeast-3.amazonaws.com<br>rolesanywhere.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | rolesanywhere.ap-northeast-2.amazonaws.com<br>rolesanywhere.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | rolesanywhere.ap-southeast-1.amazonaws.com<br>rolesanywhere.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | rolesanywhere.ap-southeast-2.amazonaws.com<br>rolesanywhere.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | rolesanywhere.ap-east-2.amazonaws.com<br>rolesanywhere.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | rolesanywhere.ap-southeast-7.amazonaws.com<br>rolesanywhere.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | rolesanywhere.ap-northeast-1.amazonaws.com<br>rolesanywhere.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | rolesanywhere.ca-central-1.amazonaws.com<br>rolesanywhere.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | rolesanywhere.ca-west-1.amazonaws.com<br>rolesanywhere.ca-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | rolesanywhere.eu-central-1.amazonaws.com<br>rolesanywhere.eu-central-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | rolesanywhere.eu-west-1.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | rolesanywhere.eu-west-2.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | rolesanywhere.eu-south-1.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | rolesanywhere.eu-west-3.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | rolesanywhere.eu-south-2.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | rolesanywhere.eu-north-1.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | rolesanywhere.eu-central-2.amazonaws.com | HTTPS | |
| | | rolesanywhere.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | rolesanywhere.il-central-1.amazonaws.com | HTTPS | |
| | | rolesanywhere.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | rolesanywhere.mx-central-1.amazonaws.com | HTTPS | |
| | | rolesanywhere.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | rolesanywhere.me-south-1.amazonaws.com<br>rolesanywhere.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | rolesanywhere.me-central-1.amazonaws.com<br>rolesanywhere.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | rolesanywhere.sa-east-1.amazonaws.com<br>rolesanywhere.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | rolesanywhere.us-gov-east-1.amazonaws.com<br>rolesanywhere.us-gov-east-1.api.aws<br>rolesanywhere-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | rolesanywhere.us-gov-west-1.amazonaws.com<br>rolesanywhere-fips.us-gov-west-1.amazonaws.com<br>rolesanywhere.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

| Resource | Description | Default value | Adjustable |
|---|---|---|---|
| Combined rate of trust anchor requests | The maximum transactions per second for ListTrustAnchors, CreateTru | 1 per second | Yes |

| Resource | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| | stAnchor, GetTrustAnchor, UpdateTrustAnchor, DeleteTrustAnchor, EnableTrustAnchor, and DisableTrustAnchor requests combined. | | |
| Combined rate of profile requests | The maximum transactions per second for ListProfiles, CreateProfile, GetProfile, UpdateProfile, DeleteProfile, EnableProfile, and DisableProfile requests combined. | 1 per second | Yes |
| Combined rate of subject requests | The maximum transactions per second for ListSubjects and GetSubject requests combined. | 1 per second | Yes |
| Combined rate of tagging requests | The maximum transactions per second for TagResource, UntagResource, and ListTagsForResource requests combined. | 1 per second | Yes |

| Resource | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Combined rate of CRL requests | The maximum transactions per second for ListCrls, GetCrl, ImportCrl, UpdateCrl, DeleteCrl , EnableCrl, and DisableCrl requests combined. | 1 per second | Yes |
| Rate of CreateSession requests | The maximum transactions per second for CreateSession requests. | 10 per second | Yes |
| Trust anchors | The maximum number of trust anchors that you can create within an account. | 50 | Yes |
| Profiles | The maximum number of profiles that you can create within an account. | 250 | Yes |
| CRLs per trust anchor | The maximum number of Certificate Revocation Lists (CRLs) that you can create per trust anchor within an account. | 2 | No |

| Resource | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Certificates per trust anchor | The maximum number of certificates that you can create per trust anchor within an account. | 2 | No |

For more information, see IAM Roles Anywhere quotas in the *IAM Roles Anywhere User Guide*.

# AWS Systems Manager Incident Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Incident Manager incidents

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | ssm-incidents.us-east-2.amazonaws.com | HTTPS |
| | | ssm-incidents-fips.us-east-2.api.aws | HTTPS |
| | | ssm-incidents-fips.us-east-2.amazonaws.com | HTTPS |
| | | ssm-incidents.us-east-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ssm-incidents.us-east-1.amazonaws.com | HTTPS |
| | | ssm-incidents-fips.us-east-1.api.aws | HTTPS |
| | | ssm-incidents-fips.us-east-1.amazonaws.com | HTTPS |
| | | ssm-incidents.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ssm-incidents.us-west-1.amazonaws.com | HTTPS |
| | | ssm-incidents-fips.us-west-1.api.aws | HTTPS |
| | | ssm-incidents-fips.us-west-1.amazonaws.com | HTTPS |
| | | ssm-incidents.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ssm-incidents.us-west-2.amazonaws.com | HTTPS |
| | | ssm-incidents-fips.us-west-2.api.aws | HTTPS |
| | | ssm-incidents-fips.us-west-2.amazonaws.com | HTTPS |
| | | ssm-incidents.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ssm-incidents.ap-south-1.amazonaws.com | HTTPS |
| | | ssm-incidents.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ssm-incidents.ap-northeast-2.amazonaws.com | HTTPS |
| | | ssm-incidents.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ssm-incidents.ap-southeast-1.amazonaws.com | HTTPS |
| | | ssm-incidents.ap-southeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | ssm-incidents.ap-southeast-2.amazonaws.com | HTTPS |
| | | ssm-incidents.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ssm-incidents.ap-northeast-1.amazonaws.com | HTTPS |
| | | ssm-incidents.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | ssm-incidents.ca-central-1.amazonaws.com | HTTPS |
| | | ssm-incidents-fips.ca-central-1.api.aws | HTTPS |
| | | ssm-incidents-fips.ca-central-1.amazonaws.com | HTTPS |
| | | ssm-incidents.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ssm-incidents.eu-central-1.amazonaws.com | HTTPS |
| | | ssm-incidents.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ssm-incidents.eu-west-1.amazonaws.com | HTTPS |
| | | ssm-incidents.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | ssm-incidents.eu-west-2.amazonaws.com | HTTPS |
| | | ssm-incidents.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | ssm-incidents.eu-west-3.amazonaws.com | HTTPS |
| | | ssm-incidents.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | ssm-incidents.eu-north-1.amazonaws.com | HTTPS |
| | | ssm-incidents.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | ssm-incidents.sa-east-1.amazonaws.com<br><br>ssm-incidents.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## Incident Manager contacts

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ssm-contacts.us-east-2.amazonaws.com<br><br>ssm-contacts.us-east-2.api.aws<br><br>ssm-contacts-fips.us-east-2.amazonaws.com<br><br>ssm-contacts-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | ssm-contacts.us-east-1.amazonaws.com<br><br>ssm-contacts.us-east-1.api.aws<br><br>ssm-contacts-fips.us-east-1.amazonaws.com<br><br>ssm-contacts-fips.us-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US West (N. California) | us-west-1 | ssm-contacts.us-west-1.amazonaws.com<br><br>ssm-contacts.us-west-1.api.aws<br><br>ssm-contacts-fips.us-west-1.amazonaws.com<br><br>ssm-contacts-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ssm-contacts.us-west-2.amazonaws.com<br><br>ssm-contacts.us-west-2.api.aws<br><br>ssm-contacts-fips.us-west-2.amazonaws.com<br><br>ssm-contacts-fips.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ssm-contacts.ap-south-1.amazonaws.com<br><br>ssm-contacts.ap-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Tokyo) | ap-northeast-1 | ssm-contacts.ap-northeast-1.amazonaws.com<br><br>ssm-contacts.ap-northeast-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ssm-contacts.ap-northeast-2.amazonaws.com<br><br>ssm-contacts.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ssm-contacts.ap-southeast-1.amazonaws.com<br><br>ssm-contacts.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ssm-contacts.ap-southeast-2.amazonaws.com<br><br>ssm-contacts.ap-southeast-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | ssm-contacts.ca-central-1.amazonaws.com<br><br>ssm-contacts.ca-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Frankfurt) | eu-central-1 | ssm-contacts.eu-central-1.amazonaws.com<br><br>ssm-contacts.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ssm-contacts.eu-west-1.amazonaws.com<br><br>ssm-contacts.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | ssm-contacts.eu-west-2.amazonaws.com<br><br>ssm-contacts.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | ssm-contacts.eu-west-3.amazonaws.com<br><br>ssm-contacts.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | ssm-contacts.eu-north-1.amazonaws.com<br><br>ssm-contacts.eu-north-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | ssm-contacts.sa-east-1.amazonaws.com<br><br>ssm-contacts.sa-east-1.api.aws | HTTPS |

# Service quotas

## Incident Manager incidents

| Name | Default | Adjustable | Description |
|---|---|---|---|
| BatchGetIncidentFindings requests per second | Each supported Region: 10 per second | Yes | The maximum number of BatchGetIncidentFindings requests per second that you can send in this account in the current region. |
| CreateReplicationSet requests per second | Each supported Region: 1 per second | Yes | The maximum number of CreateReplicationSet requests per second that you can send in this account in the current region. |
| CreateResponsePlan requests per second | Each supported Region: 5 per second | Yes | The maximum number of CreateResponsePlan requests per second that you can send in this account in the current region. |
| CreateTimelineEvent requests per second | Each supported Region: 10 per second | Yes | The maximum number of CreateTimelineEvent requests per second that you can send in this account in the current region. |
| DeleteIncidentRecord requests per second | Each supported Region: 5 per second | Yes | The maximum number of DeleteIncidentRecord requests per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | that you can send in this account in the current region. |
| DeleteReplicationSet requests per second | Each supported Region: 1 per second | Yes | The maximum number of DeleteReplicationSet requests per second that you can send in this account in the current region. |
| DeleteResourcePolicy requests per second | Each supported Region: 5 per second | Yes | The maximum number of DeleteResourcePolicy requests per second that you can send in this account in the current region. |
| DeleteResponsePlan requests per second | Each supported Region: 5 per second | Yes | The maximum number of DeleteResponsePlan requests per second that you can send in this account in the current region. |
| DeleteTimelineEvent requests per second | Each supported Region: 5 per second | Yes | The maximum number of DeleteTimelineEvent requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetIncidentRecord requests per second | Each supported Region: 10 per second | Yes | The maximum number of GetIncidentRecord requests per second that you can send in this account in the current region. |
| GetReplicationSet requests per second | Each supported Region: 10 per second | Yes | The maximum number of GetReplicationSet requests per second that you can send in this account in the current region. |
| GetResourcePolicies requests per second | Each supported Region: 10 per second | Yes | The maximum number of GetResourcePolicies requests per second that you can send in this account in the current region. |
| GetResponsePlan requests per second | Each supported Region: 10 per second | Yes | The maximum number of GetResponsePlan requests per second that you can send in this account in the current region. |
| GetTimelineEvent requests per second | Each supported Region: 30 per second | Yes | The maximum number of GetTimelineEvent requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Incidents per response plan per month | Each supported Region: 200 | No | The maximum number of incidents per response plan per month. |
| ListIncidentFindings requests per second | Each supported Region: 10 per second | Yes | The maximum number of ListIncidentFindings requests per second that you can send in this account in the current region. |
| ListIncidentRecords requests per second | Each supported Region: 10 per second | Yes | The maximum number of ListIncidentRecords requests per second that you can send in this account in the current region. |
| ListRelatedItems requests per second | Each supported Region: 10 per second | Yes | The maximum number of ListRelatedItems requests per second that you can send in this account in the current region. |
| ListReplicationSets requests per second | Each supported Region: 10 per second | Yes | The maximum number of ListReplicationSets requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ListResponsePlans requests per second | Each supported Region: 10 per second | Yes | The maximum number of ListResponsePlans requests per second that you can send in this account in the current region. |
| ListTagsForResource requests per second | Each supported Region: 10 per second | Yes | The maximum number of ListTagsForResource requests per second that you can send in this account in the current region. |
| ListTimelineEvents requests per second | Each supported Region: 30 per second | Yes | The maximum number of ListTimelineEvents requests per second that you can send in this account in the current region. |
| PutResourcePolicy requests per second | Each supported Region: 5 per second | Yes | The maximum number of PutResourcePolicy requests per second that you can send in this account in the current region. |
| Regions per replication set | Each supported Region: 3 | No | The maximum number of regions per replication set in this account. |
| Related items per incident | Each supported Region: 50 | No | The maximum number of related items per incident. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Replication sets per account | Each supported Region: 1 | No | The maximum number of replication sets in this account. |
| StartIncident requests per second | Each supported Region: 5 per second | Yes | The maximum number of StartIncident requests per second that you can send in this account in the current region. |
| TagResource requests per second | Each supported Region: 5 per second | Yes | The maximum number of TagResource requests per second that you can send in this account in the current region. |
| Timeline events per incident | Each supported Region: 1,000 | No | The maximum number of timeline events per incident. |
| UntagResource requests per second | Each supported Region: 5 per second | Yes | The maximum number of UntagResource requests per second that you can send in this account in the current region. |
| UpdateDeletionProtection requests per second | Each supported Region: 1 per second | Yes | The maximum number of UpdateDeletionProtection requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateIncidentRecord requests per second | Each supported Region: 5 per second | Yes | The maximum number of UpdateIncidentRecord requests per second that you can send in this account in the current region. |
| UpdateRelatedItems requests per second | Each supported Region: 5 per second | Yes | The maximum number of UpdateRelatedItems requests per second that you can send in this account in the current region. |
| UpdateReplicationSet requests per second | Each supported Region: 1 per second | Yes | The maximum number of UpdateReplicationSet requests per second that you can send in this account in the current region. |
| UpdateResponsePlan requests per second | Each supported Region: 5 per second | Yes | The maximum number of UpdateResponsePlan requests per second that you can send in this account in the current region. |
| UpdateTimelineEvent requests per second | Each supported Region: 5 per second | Yes | The maximum number of UpdateTimelineEvent requests per second that you can send in this account in the current region. |

## Incident Manager contacts

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AcceptPage API throttle quota | Each supported Region: 20 per second | Yes | The maximum number of AcceptPage requests per second that you can send in this account in the current region. |
| All other operations API throttle quota | Each supported Region: 1 | Yes | The maximum number of all other operation requests per second that you can send in this account in the current region. |
| Contact channels per stage | Each supported Region: 10 | Yes | The maximum number of contact channels per plan stage in this account in the current region. |
| Contacts per account | Each supported Region: 1,000 | Yes | The maximum number of contacts in this account in the current region. |
| Contacts per rotation | Each supported Region: 30 | No | The maximum number of contacts in a single rotation. |
| DescribeEngagement API throttle quota | Each supported Region: 5 per second | Yes | The maximum number of DescribeEngagement requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DescribePage API throttle quota | Each supported Region: 5 per second | Yes | The maximum number of DescribePage requests per second that you can send in this account in the current region. |
| Email engagement throttle quota | Each supported Region: 0.05 | No | The maximum number of email engagements per contact per second that the service can send in this account in the current region |
| GetContact API throttle quota | Each supported Region: 20 per second | Yes | The maximum number of GetContact requests per second that you can send in this account in the current region. |
| ListEngagements API throttle quota | Each supported Region: 2 per second | Yes | The maximum number of ListEngagements requests per second that you can send in this account in the current region. |
| ListPageReceipts API throttle quota | Each supported Region: 1 per second | Yes | The maximum number of ListPageReceipts requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListPagesByContact API throttle quota | Each supported Region: 1 per second | Yes | The maximum number of ListPagesByContact requests per second that you can send in this account in the current region. |
| ListPagesByEngagement API throttle quota | Each supported Region: 2 per second | Yes | The maximum number of ListPagesByEngagement requests per second that you can send in this account in the current region. |
| Non-mutating SSMContacts-related API throttle quota | Each supported Region: 20 | Yes | The maximum number of non-mutating SSMContacts-related API requests per second that you can send in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Non-mutating rotation-related API throttle quota | Each supported Region: 20 | Yes | The maximum number of non-mutating rotation-related API requests per second that you can send in this account in the current region. Applies individually to the GetRotation, GetRotationOverride, ListPreviewRotationShifts, ListRotationOverrides, ListPreviewRotationShifts, and ListRotationShiftsActivity commands. |
| Rotations per account | Each supported Region: 500 | Yes | The maximum number of rotations in this account in the current region. |
| Rotations per schedule | Each supported Region: 8 | No | The maximum number of rotations in a schedule. |
| SMS engagement throttle quota | Each supported Region: 0.05 | No | The maximum number of SMS engagements per contact per second that the service can send in this account in the current region |
| Stages per plan | Each supported Region: 10 | Yes | The maximum number of stages per plan in this account in the current region. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| StartEngagement API throttle quota | Each supported Region: 20 per second | Yes | The maximum number of StartEngagement requests per second that you can send in this account in the current region. |
| StopEngagement API throttle quota | Each supported Region: 10 per second | Yes | The maximum number of StopEngagement requests per second that you can send in this account in the current region. |
| Voice engagement throttle quota | Each supported Region: 0.01 | No | The maximum number of voice engagements per contact per second that the service can send in this account in the current region |

The unit for the API throttle quotas is requests per second.

## Amazon Inspector endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Amazon Inspector

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | inspector2.us-east-2.amazonaws.com<br>inspector2-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | inspector2.us-east-1.amazonaws.com<br>inspector2-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | inspector2.us-west-1.amazonaws.com<br>inspector2-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | inspector2.us-west-2.amazonaws.com<br>inspector2-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | inspector2.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | inspector2.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | inspector2.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | inspector2.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | inspector2.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | inspector2.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | inspector2.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | inspector2.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | inspector2.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | inspector2.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | inspector2.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | inspector2.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | inspector2.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | inspector2.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | inspector2.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | inspector2.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | inspector2.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | inspector2.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | inspector2.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | inspector2.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | inspector2.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | inspector2.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | inspector2.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | inspector2.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | inspector2.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | inspector2.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | inspector2.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | inspector2.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | inspector2.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | inspector2.us-gov-east-1.amazonaws.com<br>inspector2-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | inspector2.us-gov-west-1.amazonaws.com<br>inspector2-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

## Amazon Inspector Scan

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | inspector-scan.us-east-2.amazonaws.com<br>inspector-scan-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | inspector-scan.us-east-1.amazonaws.com<br>inspector-scan-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | inspector-scan.us-west-1.amazonaws.com<br>inspector-scan-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | inspector-scan.us-west-2.amazonaws.com<br>inspector-scan-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | inspector-scan.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | inspector-scan.ap-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | inspector-scan.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | inspector-scan.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | inspector-scan.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | inspector-scan.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | inspector-scan.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | inspector-scan.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | inspector-scan.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | inspector-scan.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | inspector-scan.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | inspector-scan.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | inspector-scan.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | inspector-scan.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | inspector-scan.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | inspector-scan.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | inspector-scan.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | inspector-scan.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | inspector-scan.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | inspector-scan.eu-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | inspector-scan.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | inspector-scan.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | inspector-scan.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | inspector-scan.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | inspector-scan.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | inspector-scan.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | inspector-scan.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | inspector-scan.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | inspector-scan.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | inspector-scan.us-gov-east-1.amazonaws.com<br>inspector-scan-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | inspector-scan.us-gov-west-1.amazonaws.com | HTTPS | |
| | | inspector-scan-fips.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of Suppression Rules | Each supported Region: 500 | No | The maximum number of Suppression Rules allowed per account. |

For more information, see the Amazon Inspector quotas in the *Amazon Inspector User Guide*.

# Amazon Inspector Classic endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | inspector.us-east-2.amazonaws.com<br>inspector-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | inspector.us-east-1.amazonaws.com<br>inspector-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | inspector.us-west-1.amazonaws.com<br>inspector-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | inspector.us-west-2.amazonaws.com<br>inspector-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | inspector.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | inspector.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | inspector.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | inspector.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | inspector.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | inspector.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | inspector.eu-west-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | inspector.eu-north-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | inspector.us-gov-east-1.amazonaws.com<br>inspector-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | inspector.us-gov-west-1.amazonaws.com<br>inspector-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Assessment Targets | Each supported Region: 50 | Yes | The maximum number of assessment targets that you can have at any given time per account per region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Assessment Templates | Each supported Region: 500 | Yes | The maximum number of assessment templates that you can have at any given time per account per region. |
| Assessment runs | Each supported Region: 50,000 | Yes | The maximum number of assessment runs that you can create per account per region. You can have multiple assessment runs happening at the same time as long as the assessment targets used for these runs do not contain overlapping EC2 instances. |
| Instances in running assessments | Each supported Region: 500 | Yes | The maximum number of EC2 instances that can be included across all running assessments per account per region. |

For more information, see the Amazon Inspector Classic quotas in the *Amazon Inspector User Guide*.

## AWS IoT Core endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The following sections describe the service endpoints for AWS IoT Core.

> **Note**
>
> You can use these endpoints to perform the operations in the AWS IoT API Reference. The endpoints in the following sections are different from the device endpoints, which provide devices an MQTT publish/subscribe interface and a subset of the API operations. For more information about the data, credential access, and job management endpoints used by devices, see AWS IoT device endpoints.
> For information about connecting to and using the AWS IoT endpoints, see Connecting devices to AWS IoT in the *AWS IoT Developer Guide*.

**Topics**

- AWS IoT Core - control plane endpoints
- AWS IoT Core - data plane endpoints
- AWS IoT Core - credential provider endpoints
- AWS IoT FIPS endpoints

### AWS IoT Core - control plane endpoints

The following table contains AWS Region-specific endpoints for AWS IoT Core - control plane operations. For information about the operations supported by the AWS IoT Core - control plane endpoints, see AWS IoT operations in the *AWS IoT API Reference*.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | iot.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | iot-fips.us-east-2.api.aws | HTTPS | |
| | | iot-fips.us-east-2.amazonaws.com | HTTPS | |
| | | iot.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | iot.us-east-1.amazonaws.com | HTTPS | |
| | | iot-fips.us-east-1.api.aws | HTTPS | |
| | | iot-fips.us-east-1.amazonaws.com | HTTPS | |
| | | iot.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | iot.us-west-1.amazonaws.com | HTTPS | |
| | | iot-fips.us-west-1.api.aws | HTTPS | |
| | | iot-fips.us-west-1.amazonaws.com | HTTPS | |
| | | iot.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | iot.us-west-2.amazonaws.com | HTTPS | |
| | | iot-fips.us-west-2.api.aws | HTTPS | |
| | | iot-fips.us-west-2.amazonaws.com | HTTPS | |
| | | iot.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | iot.ap-east-1.amazonaws.com | HTTPS | |
| | | iot.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | iot.ap-southeast-5.amazonaws.com | HTTPS | |
| | | iot.ap-southeast-5.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | iot.ap-south-1.amazonaws.com | HTTPS | |
| | | iot.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | iot.ap-northeast-2.amazonaws.com | HTTPS | |
| | | iot.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | iot.ap-southeast-1.amazonaws.com | HTTPS | |
| | | iot.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | iot.ap-southeast-2.amazonaws.com | HTTPS | |
| | | iot.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | iot.ap-northeast-1.amazonaws.com | HTTPS | |
| | | iot.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | iot.ca-central-1.amazonaws.com | HTTPS | |
| | | iot-fips.ca-central-1.api.aws | HTTPS | |
| | | iot-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | iot.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | iot.eu-central-1.amazonaws.com | HTTPS | |
| | | iot.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | iot.eu-west-1.amazonaws.com | HTTPS | |
| | | iot.eu-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | iot.eu-west-2.amazonaws.com | HTTPS | |
| | | iot.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | iot.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | iot.eu-west-3.amazonaws.com | HTTPS | |
| | | iot.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | iot.eu-south-2.amazonaws.com | HTTPS | |
| | | iot.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | iot.eu-north-1.amazonaws.com | HTTPS | |
| | | iot.eu-north-1.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | iot.il-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | iot.me-south-1.amazonaws.com | HTTPS | |
| | | iot.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | iot.me-central-1.amazonaws.com | HTTPS | |
| | | iot.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | iot.sa-east-1.amazonaws.com | HTTPS | |
| | | iot.sa-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | iot.us-gov-east-1.amazonaws.com | HTTPS |
| | | iot-fips.us-gov-east-1.api.aws | HTTPS |
| | | iot-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | iot.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | iot.us-gov-west-1.amazonaws.com | HTTPS |
| | | iot-fips.us-gov-west-1.api.aws | HTTPS |
| | | iot-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | iot.us-gov-west-1.api.aws | HTTPS |

### AWS IoT Core - data plane endpoints

The AWS IoT Core - data plane endpoints are specific to each AWS account and AWS Region. To find the AWS IoT Core - data plane endpoint for your AWS account and AWS Region, use the **describe-endpoint** CLI command shown here, or the `DescribeEndpoint` REST API.

```
aws iot describe-endpoint --endpoint-type iot:Data-ATS
```

This command returns your data plane API endpoint in the following format:

```
account-specific-prefix-ats.iot.aws-region.amazonaws.com
```

For information about the actions supported by the AWS IoT Core - data plane endpoints, see AWS IoT data plane operations in the *AWS IoT API Reference*.

The following table contains generic representations of the AWS account-specific endpoints for each AWS Region that AWS IoT Core supports. In the **Endpoint** column, the *`account-specific-prefix`* from your account-specific endpoint replaces `data` shown in the generic endpoint representation.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | data-ats.iot.us-east-2.amazonaws.com<br>data.iot-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | data-ats.iot.us-east-1.amazonaws.com<br>data.iot-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | data-ats.iot.us-west-1.amazonaws.com<br>data.iot-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | data-ats.iot.us-west-2.amazonaws.com<br>data.iot-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | data-ats.iot.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | data-ats.iot.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | data-ats.iot.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | data-ats.iot.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | data-ats.iot.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | data-ats.iot.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | data-ats.iot.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | data-ats.iot.ca-central-1.amazonaws.com<br>data.iot-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | data-ats.iot.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | data-ats.iot.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | data-ats.iot.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | data-ats.iot.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | data-ats.iot.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | data-ats.iot.eu-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | data-ats.iot.eu-north-1.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | data-ats.iot.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | data-ats.iot.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | data-ats.iot.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | data-ats.iot.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | data-ats.iot.us-gov-east-1.amazonaws.com | HTTPS |
| | | data.iot-fips.us-gov-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | data-ats.iot.us-gov-west-1.amazonaws.com | HTTPS |
| | | data.iot-fips.us-gov-west-1.amazonaws.com | HTTPS |

## AWS IoT Core - credential provider endpoints

**Note**

If you are an existing user of AWS IoT Core credential provider and your endpoint was previously created, the endpoint likely only supports IPv4 address by default. For dual-

stack support (to support both IPv4 and IPv6 connectivity) for your credential provider endpoint, contact AWS Support.
New credential provider endpoints support both IPv4 and IPv6 (dual-stack) by default. To find out whether your endpoint currently supports IPv6, you can run the following command:

- **For Linux** – `dig +short AAAA account-specific-prefix.credentials.iot.aws-region.amazonaws.com`
- **For Windows** – `nslookup -type=AAAA account-specific-prefix.credentials.iot.aws-region.amazonaws.com`

If this command returns no results, your endpoint does not currently support IPv6.

The AWS IoT Core credential provider endpoints are specific to each AWS account and AWS Region. To find the credential provider endpoint for your AWS account and AWS Region, use the **describe-endpoint** CLI command shown here, or the `DescribeEndpoint` REST API.

```
aws iot describe-endpoint --endpoint-type iot:CredentialProvider
```

This command returns your credential provider API endpoint in the following format:

```
account-specific-prefix.credentials.iot.aws-region.amazonaws.com
```

The following table contains generic representations of the AWS account-specific endpoints for each AWS Region that AWS IoT Core supports. In the **Endpoint** column, the *account-specific-prefix* from your account-specific endpoint replaces `prefix` shown in the generic endpoint representation.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | *prefix*.credentials.iot.us-east-2.amazonaws.com<br><br>data.credentials.iot-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | *prefix*.credentials.iot.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | data.credentials.iot-fips.us-east-1.amazonaws.com | |
| US West (N. California) | us-west-1 | *prefix*.credentials.iot.us-west-1.amazonaws.com<br><br>data.credentials.iot-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | *prefix*.credentials.iot.us-west-2.amazonaws.com<br><br>data.credentials.iot-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | *prefix*.credentials.iot.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | *prefix*.credentials.iot.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | *prefix*.credentials.iot.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | *prefix*.credentials.iot.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | *prefix*.credentials.iot.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*.credentials.iot.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*.credentials.iot.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | *prefix*.credentials.iot.ca-central-1.amazonaws.com<br><br>data.credentials.iot-fips.ca-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | *prefix*.credentials.iot.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | *prefix*.credentials.iot.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | *prefix*.credentials.iot.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | *prefix*.credentials.iot.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | *prefix*.credentials.iot.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | *prefix*.credentials.iot.eu-north-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | *prefix*.credentials.iot.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | *prefix*.credentials.iot.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | *prefix*.credentials.iot.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | *prefix*.credentials.iot.us-gov-east-1.amazonaws.com<br><br>data.credentials.iot-fips.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-West) | us-gov-west-1 | *prefix*.credentials.iot.us-gov-west-1.amazonaws.com<br><br>data.credentials.iot-fips.us-gov-west-1.amazonaws.com | HTTPS |

### AWS IoT FIPS endpoints

AWS IoT provides endpoints that support the Federal Information Processing Standard (FIPS) 140-2. Choose the appropriate FIPS compliant endpoint to access AWS IoT features in your AWS Region from FIPS Endpoints by Service. For more information about the FIPS endpoints provided by AWS IoT, see Connecting to AWS IoT FIPS endpoints.

## Service quotas

**Contents**


- AWS IoT Core thing resource limits and quotas
- AWS IoT Core thing group resource limits and quotas
- AWS IoT Core bulk thing registration limits and quotas
- AWS IoT Core billing group restrictions
- AWS IoT Core rules engine limits and quotas
- AWS IoT Core API throttling limits
- AWS IoT Core Device Shadow service limits and quotas
- AWS IoT Core Fleet Provisioning limits and quotas
- AWS IoT Core message broker and protocol limits and quotas
- AWS IoT Core credential provider limits and quotas
- AWS IoT Core security and identity limits and quotas
- MQTT-based File Delivery
- AWS IoT Core Device Advisor limits and quotas

**Note**

For the limits and quotas for the AWS IoT Core device location feature, see AWS IoT Core device location Service quotas.

## AWS IoT Core thing resource limits and quotas

### AWS IoT Core thing limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Maximum number of propagating attributes | The maximum number of propagating attributes you can add. When creating or updating a thing type, you can add propagating attributes to the published MQTT 5 messages. A propagating attribute is a key-value pair that describes aspects of an IoT resource. | 20 | No |
| Maximum number of thing attributes for a thing with a thing type | Maximum number of thing attributes for a thing with a thing type. Thing types are optional and make it easier to discover things. Things with a thing type can have up to 50 attributes. | 50 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum number of thing attributes for a thing without a thing type | Maximum number of thing attributes for a thing without a thing type. Things without a thing type can have up to three attributes. | 3 | No |
| Maximum thing name size | Maximum size of a thing name, which is 128 bytes of UTF-8 encoded characters. | 128 Bytes | No |
| Number of thing types that can be associated with a thing | Number of thing types that can be associated with a thing, which can be zero or one. Thing types are optional and their use makes it easier to discover things. | 1 | No |
| Size of thing attributes per thing | The size of thing attributes per thing, which is 47 kilobytes. Thing attributes are optional name-value pairs that store information about the thing, which makes their use easier to discover things. | 47 Kilobytes | Yes |

**Note**

**Thing types**

The number of thing types that can be defined in an AWS account is not limited.Thing types allow you to store description and configuration information that is common to all things associated with the same thing type.

## AWS IoT Core thing group resource limits and quotas

### AWS IoT Core thing group limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Maximum depth of a thing group hierarchy | The maximum depth of a hierarchy of thing groups. When you build a hierarchy of groups, the policy attached to the parent group is inherited by its child group, and by all the things in the group and its child groups. This makes it easier to manage permissio ns for large number of things. | 7 | No |
| Maximum number of attributes associated with a thing group | Maximum number of attributes associate d with a thing group. Attributes are name-value pairs you can use to store information about a | 50 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| | group. You can add, delete, or update the attributes of a group. | | |
| Maximum number of direct child groups | The maximum number of direct child groups that a thing group can have in a thing group hierarchy. | 100 | No |
| Maximum number of dynamic groups | Maximum number of dynamic groups. | 100 | No |
| Maximum number of thing groups a thing can belong to | A thing can be added to a maximum of 10 thing groups. But you cannot add a thing to more than one group in the same hierarchy . This means that a thing cannot be added to two groups that share a common parent. | 10 | No |
| Maximum size of a thing group attribute name, in chars | Maximum size of a thing group attribute name, in chars. | 128 | No |

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Maximum size of a thing group attribute value, in chars | Maximum size of a thing group attribute value, in chars. | 800 | No |
| Maximum thing group name size | Maximum thing group name size. | 128 Bytes | No |

**Note**

**Thing group assignment**

The maximum number of things that can be assigned to a thing group is not limited.

## AWS IoT Core bulk thing registration limits and quotas

### AWS IoT Core bulk thing registration

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Allowed registration tasks | For any given AWS account, only one bulk registration task can run at a time. | 1 | No |
| Data retention policy | After the bulk registration task (which can be long lived) is complete, data related to bulk | 2592000 Seconds | No |

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| | thing registration is permanently deleted after 30 days. | | |
| Maximum line length | Each line in an Amazon S3 input JSON file can't exceed 256K in length. | 256000 | No |
| Registration task termination | Any pending or incomplete bulk registration tasks are terminated after 30 days. | 2592000 Seconds | No |

For more information about the JSON file used for bulk registration, see Amazon S3 input JSON file.

## AWS IoT Core billing group restrictions

- A thing can belong to exactly one billing group.
- Unlike thing groups, billing groups cannot be organized into hierarchies.
- For its usage to be registered for tagging or billing purposes, a device must:
  - Be registered as a thing in AWS IoT Core.
  - Communicate with AWS IoT Core using MQTT only.
  - Authenticate with AWS IoT Core using only its thing name as the client ID.
  - Use an X.509 certificate or Amazon Cognito Identity to authenticate.

  For more information, see Managing Devices with AWS IoT, Authentication, and Device Provisioning. You can use the AttachThingPrincipal API operation to attach a certificate or other credential to a thing.
- The maximum number of billing groups per AWS account is 20,000.

## AWS IoT Core rules engine limits and quotas

This section describes the limits and quotas of the AWS IoT Core rules engine.

### AWS IoT Core rules engine

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum number of actions per rule | The maximum number of entries in the rule's actions property. | 10 | 10 | No |
| Maximum number of rules per AWS account | The maximum number of rules that can be defined in a single AWS account. | 1000 | 1000 | Yes |
| Rule evaluations per second per AWS account | The maximum number of rules that can be evaluated per second per AWS account. This quota includes rule evaluations that result from inbound Basic Ingest messages. | 20000 | 2000 | Yes |
| | The maximum size of a rule payload. A rule | 128 Kilobytes | 128 Kilobytes | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Rule payload size | payload includes the rule SQL, action definitions, error action definitions, rule SQL version, and description encoded in JSON using UTF-8 character encoding. Creation or modification of a rule will be rejected if the rule payload exceeds the maximum size. | | | |

*Select AWS Regions: Europe (Stockholm), Middle East (Bahrain), Middle East (UAE), Europe (Paris), Asia Pacific (Hong Kong), AWS GovCloud (US-East), AWS GovCloud (US-West), US West (N. California), Canada (Central), China (Ningxia), Asia Pacific (Malaysia), Europe (Spain)

**AWS IoT Core rules engine HTTP actions limits and quotas**

**AWS IoT Core HTTP action**

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| HTTP Action: Maximum length | Maximum length of an endpoint URL for topic rule HTTP Action. | 2 Kilobytes | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| of an endpoint URL | | | |
| HTTP Action: Maximum number of headers per action | Maximum number of headers per HTTP action. When specifying the list of headers to include in the HTTP request, it must contain a header key and a header value. To learn more, see https://docs.aws.amazon.com/iot/latest/developerguide/https-rule-action.html. | 100 | No |
| HTTP Action: Maximum size of a header key | Maximum size of a header key for topic rule HTTP action. The header file for a HTTP request includes this header key and a header value. | 256 Bytes | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| HTTP Action: Maximum topic rule destinations per AWS account | Maximum number of topic rule destinations per AWS account for topic rule HTTPS action. You must confirm and enable HTTPS endpoints before the rules engine can use them. For more information, see https://docs.aws.amazon.com/iot/latest/developerguide/rule-destination.html. | 1000 | No |
| HTTP Action: Request timeout | Request timeout for topic rule HTTP action. The AWS IoT rules engine retries the HTTPS action until the total time to complete a request exceeds the timeout quota. | 3000 Milliseconds | No |

| Resource | Value | Adjustable |
| --- | --- | --- |
| TCP ports used for HTTP actions | 443, 8443 | No |

### AWS IoT Core rules engine Apache Kafka actions limits and quotas

| Resource | Limits |
|---|---|
| Bootstrap server ports | 9000-9100 |
| Kerberos key distribution center (KDC) | 88 |

### AWS IoT Core rules engine VPC actions limits and quotas

| Resource | Quota |
|---|---|
| Maximum number of VPC destinations | 5 per account per Region |

## AWS IoT Core API throttling limits

This table describes the maximum number of transactions per second (TPS) that can be made to each of these AWS IoT Core API actions.

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| AcceptCertificateTransfer API TPS | The maximum number of transactions per second (TPS) that can be made for the AcceptCertificateTransfer API. | 10 | 10 | Yes |
| AddThingToBillingG | The maximum number of transactions | 60 | 60 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| roup API TPS | per second (TPS) that can be made for the AddThingToBillingGroup API. | | | |
| AddThingToThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the AddThingToThingGroup API. | 100 | 60 | Yes |
| AttachPolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the AttachPolicy API. | 15 | 15 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| AttachPrincipalPolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the AttachPrincipalPolicy API. | 15 | 15 | Yes |
| AttachThingPrincipal API TPS | The maximum number of transactions per second (TPS) that can be made for the AttachThingPrincipal API. | 100 | 50 | Yes |
| CancelCertificateTransfer API TPS | The maximum number of transactions per second (TPS) that can be made for the CancelCertificateTransfer API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ClearDefaultAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the ClearDefaultAuthorizer API. | 10 | 10 | Yes |
| CreateAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateAuthorizer API. | 10 | 10 | No |
| CreateBillingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateBillingGroup API. | 25 | 25 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| CreateCertificateFromCsr API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateCertificateFromCsr API. | 15 | 15 | Yes |
| CreateCertificateProvider API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateCertificateProvider API. | 1 | 1 | No |
| CreateDomainConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateDomainConfiguration API. | 1 | 1 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| CreateDynamicThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateDynamicThingGroup API. | 5 | 5 | Yes |
| CreateKeysAndCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateKeysAndCertificate API. | 10 | 10 | Yes |
| CreatePolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the CreatePolicy API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| CreatePolicyVersion API TPS | The maximum number of transactions per second (TPS) that can be made for the CreatePolicyVersion API. | 10 | 10 | Yes |
| CreateProvisioningClaim API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateProvisioningClaim API. | 10 | 10 | Yes |
| CreateProvisioningTemplate API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateProvisioning Template API. | 10 | 10 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| CreateProvisioningTemplateVersion API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateProvisioning TemplateVersion API. | 10 | 10 | No |
| CreateRoleAlias API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateRoleAlias API. | 10 | 10 | No |
| CreateThing API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateThing API. | 100 | 50 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| CreateThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateThingGroup API. | 25 | 25 | Yes |
| CreateThingType API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateThingType API. | 15 | 15 | Yes |
| CreateTopicRule API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateTopicRule API. | 5 | 5 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| CreateTopicRuleDestination API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateTopicRuleDestination API. | 5 | 5 | No |
| DeleteAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteAuthorizer API. | 10 | 10 | No |
| DeleteBillingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteBillingGroup API. | 15 | 15 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeleteCACertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteCACertificate API. | 10 | 10 | Yes |
| DeleteCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteCertificate API. | 10 | 10 | Yes |
| DeleteCertificateProvider API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteCertificateProvider API. | 1 | 1 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeleteDomainConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteDomainConfiguration API. | 10 | 10 | No |
| DeleteDynamicThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteDynamicThingGroup API. | 5 | 5 | Yes |
| DeletePolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the DeletePolicy API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeletePolicyVersion API TPS | The maximum number of transactions per second (TPS) that can be made for the DeletePolicyVersion API. | 10 | 10 | Yes |
| DeleteProvisioningTemplate API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteProvisioningTemplate API. | 10 | 10 | Yes |
| DeleteProvisioningTemplateVersion API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteProvisioningTemplateVersion API. | 10 | 10 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeleteRegistrationCode API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteRegistrationCode API. | 10 | 10 | Yes |
| DeleteRoleAlias API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteRoleAlias API. | 10 | 10 | No |
| DeleteThing API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteThing API. | 100 | 50 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeleteThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteThingGroup API. | 15 | 15 | Yes |
| DeleteThingType API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteThingType API. | 15 | 15 | Yes |
| DeleteTopicRule API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteTopicRule API. | 20 | 5 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeleteTopicRuleDestination API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteTopicRuleDestination API. | 5 | 5 | No |
| DeleteV2LoggingLevel API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteV2LoggingLevel API. | 2 | 2 | No |
| DeprecateThingType API TPS | The maximum number of transactions per second (TPS) that can be made for the DeprecateThingType API. | 15 | 15 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DescribeAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeAuthorizer API. | 10 | 10 | Yes |
| DescribeBillingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeBillingGroup API. | 100 | 100 | Yes |
| DescribeCACertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeCACertificate API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DescribeCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeCertificate API. | 10 | 10 | Yes |
| DescribeCertificateProvider API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeCertificateProvider API. | 10 | 10 | No |
| DescribeDefaultAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeDefaultAuthorizer API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DescribeDomainConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeDomainConfiguration API. | 10 | 10 | Yes |
| DescribeEncryptionConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeEncryptionConfiguration API. | 10 | 10 | No |
| DescribeEndpoint API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeEndpoint API. | 10 | 10 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DescribeEventConfigurations API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeEventConfigurations API. | 10 | 10 | Yes |
| DescribeProvisioningTemplate API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeProvisioningTemplate API. | 10 | 10 | Yes |
| DescribeProvisioningTemplateVersion API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeProvisioningTemplateVersion API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DescribeRoleAlias API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeRoleAlias API. | 10 | 10 | Yes |
| DescribeThing API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeThing API. | 350 | 350 | Yes |
| DescribeThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeThingGroup API. | 100 | 100 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DescribeThingRegistrationTask API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeThingRegistrationTask API. | 10 | 10 | Yes |
| DescribeThingType API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeThingType API. | 100 | 50 | Yes |
| DetachPolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the DetachPolicy API. | 15 | 15 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DetachPrincipalPolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the DetachPrincipalPolicy API. | 15 | 15 | Yes |
| DetachThingPrincipal API TPS | The maximum number of transactions per second (TPS) that can be made for the DetachThingPrincipal API. | 100 | 50 | Yes |
| DisableTopicRule API TPS | The maximum number of transactions per second (TPS) that can be made for the DisableTopicRule API. | 5 | 5 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| EnableTopicRule API TPS | The maximum number of transactions per second (TPS) that can be made for the EnableTopicRule API. | 5 | 5 | No |
| GetEffectivePolicies API TPS | The maximum number of transactions per second (TPS) that can be made for the GetEffectivePolicies API. | 5 | 5 | No |
| GetLoggingOptions API TPS | The maximum number of transactions per second (TPS) that can be made for the GetLoggingOptions API. | 2 | 2 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| GetPolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the GetPolicy API. | 10 | 10 | Yes |
| GetPolicy Version API TPS | The maximum number of transactions per second (TPS) that can be made for the GetPolicyVersion API. | 15 | 15 | Yes |
| GetRegist rationCode API TPS | The maximum number of transactions per second (TPS) that can be made for the GetRegist rationCode API. | 10 | 10 | Yes |
| GetRetain edMessage API TPS | The maximum number of transactions per second that can be made for the GetRetain edMessage API. | 500 | 50 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| GetTopicRule API TPS | The maximum number of transactions per second (TPS) that can be made for the GetTopicRule API. | 200 | 20 | No |
| GetTopicRuleDestination API TPS | The maximum number of transactions per second (TPS) that can be made for the GetTopicRuleDestination API. | 50 | 5 | No |
| GetV2LoggingOptions API TPS | The maximum number of transactions per second (TPS) that can be made for the GetV2LoggingOptions API. | 2 | 2 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListAttachedPolicies API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAttachedPolicies API. | 15 | 15 | Yes |
| ListAuthorizers API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAuthorizers API. | 10 | 10 | Yes |
| ListBillingGroups API TPS | The maximum number of transactions per second (TPS) that can be made for the ListBillingGroups API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListCACertificates API TPS | The maximum number of transactions per second (TPS) that can be made for the ListCACertificates API. | 10 | 10 | Yes |
| ListCertificateProviders API TPS | The maximum number of transactions per second (TPS) that can be made for the ListCertificateProviders API. | 10 | 10 | No |
| ListCertificates API TPS | The maximum number of transactions per second (TPS) that can be made for the ListCertificates API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListCertificatesByCA API TPS | The maximum number of transactions per second (TPS) that can be made for the ListCertificatesByCA API. | 10 | 10 | Yes |
| ListDomainConfigurations API TPS | The maximum number of transactions per second (TPS) that can be made for the ListDomainConfigurations API. | 10 | 10 | Yes |
| ListOutgoingCertificates API TPS | The maximum number of transactions per second (TPS) that can be made for the ListOutgoingCertificates API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListPolicies API TPS | The maximum number of transactions per second (TPS) that can be made for the ListPolicies API. | 10 | 10 | Yes |
| ListPolicyPrincipals API TPS | The maximum number of transactions per second (TPS) that can be made for the ListPolicyPrincipals API. | 10 | 10 | Yes |
| ListPolicyVersions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListPolicyVersions API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListPrincipalPolicies API TPS | The maximum number of transactions per second (TPS) that can be made for the ListPrincipalPolicies API. | 15 | 15 | Yes |
| ListPrincipalThings API TPS | The maximum number of transactions per second (TPS) that can be made for the ListPrincipalThings API. | 10 | 10 | Yes |
| ListPrincipalThingsV2 API TPS | The maximum number of transactions per second (TPS) that can be made for the ListPrincipalThingsV2 API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListProvisioningTemplateVersions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListProvisioningTemplateVersions API. | 10 | 10 | Yes |
| ListProvisioningTemplates API TPS | The maximum number of transactions per second (TPS) that can be made for the ListProvisioningTemplates API. | 10 | 10 | Yes |
| ListRetainedMessages API TPS | The maximum number of transactions per second that can be made for the ListRetainedMessages API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListRoleAliases API TPS | The maximum number of transactions per second (TPS) that can be made for the ListRoleAliases API. | 10 | 10 | Yes |
| ListTagsForResource API TPS | The maximum number of transactions per second (TPS) that can be made for the ListTagsForResource API. | 10 | 10 | Yes |
| ListTargetsForPolicy API TPS | The maximum number of transactions per second (TPS) that can be made for the ListTargetsForPolicy API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListThingGroups API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThingGroups API. | 10 | 10 | Yes |
| ListThingGroupsForThing API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThingGroupsForThing API. | 100 | 50 | Yes |
| ListThingPrincipals API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThingPrincipals API. | 20 | 20 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListThingPrincipalsV2 API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThing PrincipalsV2 API. | 20 | 20 | Yes |
| ListThingRegistrationTaskReports API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThing Registrat ionTaskReports API. | 10 | 10 | Yes |
| ListThingRegistrationTasks API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThing Registrat ionTasks API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListThingTypes API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThingTypes API. | 10 | 10 | Yes |
| ListThings API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThings API. | 10 | 10 | Yes |
| ListThingsInBillingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThingsInBillingGroup API. | 25 | 25 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListThingsInThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the ListThingsInThingGroup API. | 25 | 25 | Yes |
| ListTopicRuleDestinations API TPS | The maximum number of transactions per second (TPS) that can be made for the ListTopicRuleDestinations API. | 1 | 1 | No |
| ListTopicRules API TPS | The maximum number of transactions per second (TPS) that can be made for the ListTopicRules API. | 2 | 2 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListV2LoggingLevels API TPS | The maximum number of transactions per second (TPS) that can be made for the ListV2LoggingLevels API. | 2 | 2 | No |
| Maximum number of resource-specific logging configurations per AWS account | The maximum number of resource-specific logging configurations that can be defined in a single AWS account. | 1000 | 1000 | No |
| RegisterCACertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the RegisterCACertificate API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| RegisterCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the RegisterCertificate API. | 10 | 10 | Yes |
| RegisterCertificateWithoutCA API TPS | The maximum number of transactions per second (TPS) that can be made for the RegisterCertificateWithoutCA API. | 10 | 10 | Yes |
| RegisterThing API TPS | The maximum number of transactions per second (TPS) that can be made for the RegisterThing API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| RejectCertificateTransfer API TPS | The maximum number of transactions per second (TPS) that can be made for the RejectCertificateTransfer API. | 10 | 10 | Yes |
| RemoveThingFromBillingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the RemoveThingFromBillingGroup API. | 30 | 30 | Yes |
| RemoveThingFromThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the RemoveThingFromThingGroup API. | 100 | 50 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ReplaceTopicRule API TPS | The maximum number of transactions per second (TPS) that can be made for the ReplaceTopicRule API. | 5 | 5 | No |
| SetDefaultAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the SetDefaultAuthorizer API. | 10 | 10 | Yes |
| SetDefaultPolicyVersion API TPS | The maximum number of transactions per second (TPS) that can be made for the SetDefaultPolicyVersion API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| SetLoggingOptions API TPS | The maximum number of transactions per second (TPS) that can be made for the SetLoggingOptions API. | 2 | 2 | No |
| SetV2LoggingLevel API TPS | The maximum number of transactions per second (TPS) that can be made for the SetV2LoggingLevel API. | 2 | 2 | No |
| SetV2LoggingOptions API TPS | The maximum number of transactions per second (TPS) that can be made for the SetV2LoggingOptions API. | 2 | 2 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| StartThingRegistrationTask API TPS | The maximum number of transactions per second (TPS) that can be made for the StartThingRegistrationTask API. | 10 | 10 | Yes |
| StopThingRegistrationTask API TPS | The maximum number of transactions per second (TPS) that can be made for the StopThingRegistrationTask API. | 10 | 10 | Yes |
| TagResource API TPS | The maximum number of transactions per second (TPS) that can be made for the TagResource API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| TestAuthorization API TPS | The maximum number of transactions per second (TPS) that can be made for the TestAuthorization API. | 10 | 10 | No |
| TestInvokeAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the TestInvokeAuthorizer API. | 10 | 10 | No |
| TransferCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the TransferCertificate API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| UntagResource API TPS | The maximum number of transactions per second (TPS) that can be made for the UntagResource API. | 10 | 10 | Yes |
| UpdateAuthorizer API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateAuthorizer API. | 10 | 10 | Yes |
| UpdateBillingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateBillingGroup API. | 15 | 15 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| UpdateCACertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateCACertificate API. | 10 | 10 | Yes |
| UpdateCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateCertificate API. | 10 | 10 | Yes |
| UpdateCertificateMode API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateCertificateMode API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| UpdateCertificateProvider API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateCertificateProvider API. | 1 | 1 | No |
| UpdateDomainConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateDomainConfiguration API. | 10 | 10 | Yes |
| UpdateDynamicThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateDynamicThingGroup API. | 5 | 5 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| UpdateEncryptionConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateEncryptionConfiguration API. | 1 | 1 | No |
| UpdateEventConfigurations API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateEventConfigurations API. | 10 | 10 | Yes |
| UpdateProvisioningTemplate API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateProvisioningTemplate API. | 10 | 10 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| UpdateRoleAlias API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateRoleAlias API. | 10 | 10 | Yes |
| UpdateThing API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateThing API. | 100 | 50 | Yes |
| UpdateThingGroup API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateThingGroup API. | 15 | 15 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| UpdateThingGroupsForThing API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateThingGroupsForThing API. | 10 | 10 | Yes |
| UpdateThingType API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateThingType API. | 15 | 15 | Yes |
| UpdateTopicRuleDestination API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateTopicRuleDestination API. | 5 | 5 | No |

# AWS IoT Core Device Shadow service limits and quotas

## AWS IoT Core Device Shadow actions

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Device Shadow API requests/second per account | Number of device shadow API requests per second per account. This value is adjustable and subject to per-account quotas, depending on the region. | 4000 | 400 | Yes |
| Maximum depth of JSON device state documents | The maximum number of levels in the desired or reported section of the JSON device state document is 8. | 8 | 8 | No |
| Maximum number of in-flight, unacknowledged messages per thing | The Device Shadow service supports up to 10 in-flight unacknowledged messages per thing on a single connection. When this quota is reached, all | 10 | 10 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| | new shadow requests are rejected with a 429 error code until the number of in-flight requests drop below the limit. | | | |
| Maximum shadow name size | Maximum size of a thing shadow name, which is 64 bytes of UTF-8 encoded characters. | 64 Bytes | 64 Bytes | No |
| Maximum size of a JSON state document | Each individual shadow document must be 8KB or less in size. Metadata doesn't contribute to the document size for service quotas or pricing. | 8 Kilobytes | 8 Kilobytes | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum thing name size | Maximum size of a thing name, which is 128 bytes of UTF-8 encoded characters. | 128 Bytes | 128 Bytes | No |
| Requests per second per thing | The Device Shadow service supports up to 20 requests per second per shadow. This limit applies to both HTTP and MQTT API calls for GetThingShadow, UpdateThingShadow, DeleteThingShadow, and ListNamedShadowsForThing (classic shadows only). | 20 | 20 | Yes |

*Select AWS Regions: Europe (Paris), Europe (Stockholm), Asia Pacific (Hong Kong), South America (São Paulo), Canada (Central), Middle East (Bahrain), Middle East (UAE), China (Ningxia), AWS GovCloud (US-East), AWS GovCloud (US-West) Asia Pacific (Malaysia), Europe (Spain)

The levels in the `desired` and `reported` sections of the Device Shadow's JSON state document are counted as shown here for the `desired` object.

```
"desired": {
    "one": {
        "two": {
            "three": {
                "four": {
                    "five":{
                    }
                }
            }
        }
    }
}
```

> **Note**
>
> AWS IoT Core deletes a Device Shadow document after the creating account is deleted or upon customer request. For operational purposes, AWS IoT service backups are retained for 6 months.

## AWS IoT Core Fleet Provisioning limits and quotas

Following are throttling limits for some fleet provisioning APIs per AWS account.

### AWS IoT Core fleet provisioning limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Fleet Provisioning CreateCertificateFromCsr MQTT API TPS | The maximum number of transactions per second (TPS) that can be made for the Fleet Provisioning CreateCertificateFromCsr MQTT API. | 20 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Fleet Provisioning CreateKeysAndCertificate MQTT API TPS | The maximum number of transactions per second (TPS) that can be made for the Fleet Provisioning CreateKeysAndCertificate MQTT API. | 10 | Yes |
| Fleet Provisioning RegisterThing MQTT API TPS | The maximum number of transactions per second (TPS) that can be made for the Fleet Provisioning RegisterThing MQTT API. | 10 | Yes |

Fleet provisioning also has these limits, which can't be changed.

| Resource | Description | Limit | |
| --- | --- | --- | --- |
| Versions per fleet provisioning template | The maximum number of versions that a fleet provisioning template can have. Each template version has a version ID and a creation date for devices that connect to AWS IoT using fleet provisioning. | 5 | |
| Fleet provisioning templates per customer | The maximum number of fleet provisioning | 256 | |

| Resource | Description | Limit | |
|---|---|---|---|
| | templates per customer. Use fleet provisioning templates to generate certificates and private keys for your devices to securely connect to AWS IoT. | | |
| Fleet provisioning template size | The maximum size of a fleet provisioning template in Kilobytes. Fleet provisioning templates allow you to generate certificates and private keys for your devices to securely connect to AWS IoT. | 10 Kilobytes | |

## AWS IoT Core message broker and protocol limits and quotas

**Note**

The limits listed below are per AWS Regions.

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Client ID size | Size of the client ID, which is 128 bytes of | 128 Bytes | 128 Bytes | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| | UTF-8 encoded characters. | | | |
| Connect requests per second per account | The maximum number of MQTT CONNECT requests per second per account. | 3000 | 100 | Yes |
| Connect requests per second per client ID | AWS IoT Core restricts MQTT CONNECT requests from the same accountId and clientId to 1 MQTT CONNECT operation per second. | 1 | 1 | No |
| DeleteConnection requests per second per client ID | AWS IoT Core restricts DeleteConnection requests from the same account and client ID to 1 request per second. | 1 | 1 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| DeleteConnection API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteConnection API. | 100 | 100 | Yes |
| GetConnection requests per second per client ID | AWS IoT Core restricts GetConnection requests from the same account and client ID to 1 request per second. | 1 | 1 | No |
| GetConnection API TPS | The maximum number of GetConnection API requests per second per account. | 100 | 100 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| ListSubscriptions requests per second per client ID | AWS IoT Core restricts ListSubscriptions requests from the same account and client ID to 1 request per second. | 1 | 1 | No |
| ListSubscriptions API TPS | The maximum number of ListSubscriptions API requests per second per account. | 100 | 100 | Yes |
| SendDirectMessage API TPS | The maximum number of SendDirectMessage API requests per second per account. | 20000 | 2000 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Connection inactivity (keep-alive interval) | The default keep-alive interval is used when a client requests a keep-alive interval of zero or > 1200 seconds. If a client requests a keep-alive interval < 30 seconds but more than zero, the server treats the client as though it requested a keep-alive interval of 30 seconds. | 1200 Seconds | 1200 Seconds | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Inbound publish requests per second per account | Inbound publish requests counts all messages that IoT Core processes before routing them to the clients or rules engine. Ex: A single message published on reserved topic can result in publishing 3 additional messages for shadow update, documents and delta, hence counted as 4 requests; whereas on an unreserved topic like a/b is counted as 1 request. | 20000 | 2000 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| MQTT payload size | The payload for every publish request can be no larger than 128 KB. AWS IoT Core rejects publish and connect requests larger than this size. | 128 Kilobytes | 128 Kilobytes | No |
| Maximum MQTT5 Content Type size | The maximum size for MQTT5 Content Type (UTF-8 string). | 256 Bytes | 256 Bytes | No |
| Maximum MQTT5 Correlation Data size | The maximum size for MQTT5 Correlation Data. | 8 Kilobytes | 8 Kilobytes | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum MQTT5 Topic Alias value | The maximum value of MQTT 5 Topic Alias topics that a client can specify per connection allowed. If the client attempts to publish with Topic Alias topics greater than the maximum value, the client will be disconnected. | 8 | 8 | Yes |
| Maximum MQTT5 packet size | The maximum MQTT5 packet size (variable header and payload) | 146 Kilobytes | 146 Kilobytes | No |
| Maximum Message Expiry Interval | The maximum MQTT5 lifetime of a message in seconds. | 604800 Seconds | 604800 Seconds | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum User Properties total size | The maximum total size per packet in kilobytes for MQTT5 User Property keys and values. | 8 Kilobytes | 8 Kilobytes | No |
| Maximum concurrent client connections per account | The maximum number of concurrent connections allowed per account. | 500000 | 100000 | Yes |
| Maximum inbound unacknowledged QoS 1 publish requests | AWS IoT Core restricts the number of unacknowledged inbound publish requests per MQTT client. When this quota is reached, no new publish requests are accepted from this client until a PUBACK message is returned by the server. | 100 | 100 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum number of retained messages per account | The number of stored retained messages per account.When this limit is reached, no new retained messages are stored for this account and all retained publishes with payloads greater than 0 bytes are throttled. | 500000 | 100000 | Yes |
| Maximum number of slashes in topic and topic filter | A topic in a publish or subscribe request can have no more than 7 forward slashes (/). This excludes the first 3 slashes in the mandatory segments for Basic Ingest topics ($AWS/rules/rule-name/). | 7 | 7 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum outbound unacknowledged QoS 1 publish requests | AWS IoT Core restricts the number of unacknowledged outbound publish requests per client. When this quota is reached, no new publish requests are sent to the client until the client acknowledges the publish requests. | 100 | 100 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Maximum retry interval for delivering QoS 1 messages | AWS IoT Core retries delivery of unacknowledged quality of service 1 (QoS 1) publish requests to a connected client for up to one hour for MQTT 3 subscribers and sixteen minutes for MQTT 5 subscribers, after which it drops the publish requests. | 960 Seconds | 960 Seconds | No |
| Maximum subscriptions per subscribe request | A single SUBSCRIBE request has a quota of 8 subscriptions. | 8 | 8 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Minimum MQTT5 maximum packet size | To connect to AWS IoT Core, clients can specify a maximum packet size that it supports. Clients cannot set this value to less than 512 bytes. Clients that specify values less than 512 bytes will not be allowed to connect. | 512 Bytes | 512 Bytes | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Outbound publish requests per second per account | Outbound publish requests count for every message that resulted in matching a client's subscription. For example, 2 clients are subscribed to topic filter a/b. An inbound publish request on topic a/b results in a total of 2 outbound publish requests. For topics with a large number of subscribers, message delivery is subject to increased latency. AWS IoT Core also evaluates this quota when processing inbound publish requests, | 20000 | 2000 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| | because outbound requests originate from them. If the quota is exceeded, AWS IoT Core throttles the inbound publish request and drops the message. | | | |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Persistent session expiry period | The duration for which the message broker stores MQTT persistent sessions, their queued messages, and messages in shared subscription group queues. When a session disconnects, the expiry timer starts. After expiry, the broker terminates the session and discards all queued messages. | 3600 Seconds | 3600 Seconds | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Publish requests per second per connection | AWS IoT Core restricts each client connection to a maximum number of inbound and outbound publish requests per second. This limit includes messages sent to offline persistent session and messages sent via SendDirectMessage. Publish requests that exceed that quota are discarded. | 100 | 100 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Queued messages per second per account | The maximum number of messages that AWS IoT Core can queue per second for an account. This limit applies when AWS IoT Core stores messages sent to offline persistent sessions or shared subscription groups. | 500 | 500 | Yes |
| Queued Messages per shared subscription group | Maximum approximate number of queued messages for each shared subscription group. | 500000 | 100000 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Retained message inbound publish requests per second per account | The maximum rate that AWS IoT Core can accept inbound publish requests of MQTT messages with the RETAIN flag set.This rate includes all inbound publish requests whether invoked by the HTTP or MQTT protocol. | 500 | 50 | Yes |
| Retained message inbound publish requests per second per topic | MQTT/HTTP publish requests with RETAIN flag set made to the same topic per second. | 1 | 1 | No |
| Shared Subscription groups per account | The maximum number of shared subscription groups per account. | 100 | 100 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Shared Subscriptions per group | The maximum number of subscriptions in a shared subscription group. | 100 | 100 | No |
| Subscribe and unsubscribe requests per second per group | The maximum number of subscribe and unsubscribe requests that AWS IoT Core can accept per second in a shared subscription group. | 100 | 100 | No |
| Subscriptions per account | AWS IoT Core restricts an account to a maximum number of subscriptions across all active connections. | 500000 | 100000 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Subscriptions per connection | AWS IoT Core restricts the number of subscriptions per connection. AWS IoT Core rejects subscription requests on the same connection in excess of this amount. Clients should validate the SUBACK message to ensure that their subscription requests have been successfully processed. | 50 | 50 | Yes |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Subscriptions per second per account | AWS IoT Core restricts an account to a maximum number of subscriptions per second. For example, if there are 2 MQTT SUBSCRIBE requests sent within a second, each with 3 subscriptions (topic filters), AWS IoT Core counts those as 6 subscriptions. | 3000 | 200 | Yes |
| Throughput per second per connection | Data received or sent over a client connection is processed at a maximum throughput rate. Data that exceeds the maximum throughput is delayed in processing. | 512 Kilobytes | 512 Kilobytes | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| Topic size | The topic passed to AWS IoT Core when sending a publish request can be no larger than 256 bytes of UTF-8 encoded characters. This excludes the first 3 mandatory segments for Basic Ingest topics ($AWS/rules/rule-name/). | 256 Bytes | 256 Bytes | No |
| WebSocket connection duration | The WebSocket connection lifetime is 24 hours. If the lifetime is exceeded, The WebSocket connection will be closed. | 86400 Seconds | 86400 Seconds | No |

*Select AWS Regions: Europe (Stockholm), Middle East (Bahrain), Middle East (UAE), Europe (Paris), Asia Pacific (Hong Kong), AWS GovCloud (US-East), AWS GovCloud (US-West), US West (N. California), Canada (Central), China (Ningxia), Asia Pacific (Malaysia), Europe (Spain), South America (São Paulo), Europe (Milan), Israel (Tel Aviv)

# AWS IoT Core credential provider limits and quotas

## AWS IoT Core credential limits and quotas

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| AssumeRoleWithCertificate API TPS | The maximum number of transactions per second (TPS) that can be made for the AssumeRoleWithCertificate API. | 50 | 100 | Yes |
| Maximum number of AWS IoT Core role aliases | Maximum number of AWS IoT Core role aliases registered in your AWS account. AWS IoT role alias allows connected devices to authenticate to AWS IoT using X.509 certificates and obtain short-lived AWS credentials from an IAM role that is associated | 100 | 100 | No |

| Limit display name | Description | Default value | Default value in select AWS Regions* | Adjustable |
|---|---|---|---|---|
| | with the role alias. | | | |

*Select AWS Regions: US East (N. Virginia), US West (Oregon), Europe (Ireland) Asia Pacific (Malaysia), Europe (Spain)

> **Note**
>
> **Large Region** limits apply to AWS Regions: US East (N. Virginia), US West (Oregon), and Europe (Ireland)

## AWS IoT Core security and identity limits and quotas

### AWS IoT Core security and identity limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Configurable endpoints: maximum number of domain configurations per account | Configurable endpoints: maximum number of domain configurations per account | 10 | Yes |
| Custom authentication: maximum number of active authorizers per account | Custom authentication: maximum number of authorizers that can be active in your AWS account. | 32 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Custom authentication: maximum number of authorizers per account | Custom authentication: maximum number of authorizers that can be registered to your AWS account. Authorizers have a lambda function that implements custom authentication and authorization. | 100 | No |
| Maximum number of CA certificates with the same subject field allowed per AWS account per Region | The maximum number of CA certificates with the same subject field allowed per AWS account per region. If you have more than one CA certificate with the same subject field, you must specify the CA certificate that was used to sign the device certificate being registered. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum number of device certificates that can be registered per second | The maximum number of device certificates that can be registered per second. You can select up to 15 files to register. | 15 | Yes |
| Maximum number of domain configurations per account per region | Maximum number of domain configurations per AWS account per AWS Region. | 10 | Yes |
| Maximum number of fleet provisioning template versions per template | Maximum number of fleet provisioning template versions per template. Each template version has a version ID and a creation date for devices connecting to AWS IoT using fleet previsioning. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum number of fleet provisioning templates per customer | Maximum number of fleet provisioning templates per customer. Use fleet provisioning templates to generate certificates and private keys for your devices to securely connect to AWS IoT. | 256 | No |
| Maximum number of named policy versions | The maximum number of named policy versions. A managed AWS IoT policy can have up to five versions. To update a policy, create a new policy version. If the policy has five versions, you must delete an existing version before creating a new one. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum number of policies that can be attached to a certificate or Amazon Cognito identity | The maximum number of policies that can be attached to a client certificate or an Amazon Cognito identity, which is 10. Amazon Cognito identity enables you to create temporary, limited-privilege AWS credentials for use in mobile and web applications. | 10 | No |
| Maximum number of provisioning claims that can be generated per second by trusted user | The maximum number of provisioning claims that can be generated per second by a trusted user. A trusted user can be an end user or installation technician who uses a mobile app or web application to configure the device in its deployed location. | 10 | No |
| Maximum policy document size | The maximum size of the policy document, which is 2048 characters excluding white spaces. | 2048 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum size of fleet provisioning template | Maximum size of fleet provisioning templates in Kilobytes. Fleet provisioning templates allow you to generate certificates and private keys for your devices to securely connect to AWS IoT. | 10 Kilobytes | No |

### Additional AWS IoT Core security limits

| Resource | Description | Default | Adjustable |
| --- | --- | --- | --- |
| Maximum number of domain names per server certificate | When you're providing the server certificates for AWS IoT custom domain configuration, certificates can have a maximum of four domain names. | 4 | No |
| Custom authentication: minimum connection duration (value of `DisconnectAfterInSecs`) | The Lambda function of a custom authorizer uses a `DisconnectAfterInSecs` parameter to indicate the maximum duration (in seconds) of the connection to the AWS IoT Core gateway. The connection is terminated if it exceeds this value. The minimum value for `DisconnectAfterInSecs` is 300. If the value is less than 300, a client can be disconnected by AWS IoT Core. A connection can't last less than 300 seconds (the minimum value). | 300 | No |

| Resource | Description | Default | Adjustable |
| --- | --- | --- | --- |
| Custom authentication: maximum connection duration (value of `DisconnectAfterInSecs`) | The Lambda function of a custom authorizer uses a `DisconnectAfterInSecs` parameter to indicate the maximum duration (in seconds) of the connection to the AWS IoT Core gateway. The connection is terminated if it exceeds this value. The maximum value for `DisconnectAfterInSecs` is 86,400. If the value is more than 86,400, a client can be disconnected by AWS IoT Core. A connection can't last longer than 86,400 seconds (the maximum value). This applies to all custom authentication connections (MQTT and WSS). | 86,400 | No |
| Custom authentication: minimum policy refresh rate (value of `RefreshAfterInSecs`) | The Lambda function of a custom authorizer uses a `RefreshAfterInSeconds` parameter to indicate the interval (in seconds) between policy refreshes when connected to the AWS IoT Core gateway. When this interval passes, AWS IoT Core invokes the Lambda function to allow for policy refreshes. | 300 | No |
| Custom authentication: maximum policy refresh rate (value of `RefreshAfterInSecs`) | The maximum time interval between policy refreshes when connected to the AWS IoT Core gateway, defined by the value of `RefreshAfterInSeconds`. | 86,400 | No |

# MQTT-based File Delivery

## MQTT-based File Delivery Resource Quotas

| Resource | Description | Default | Adjustable |
|---|---|---|---|
| Streams per account | The maximum number of streams per account. | 10,000* | No |
| Files per stream | The maximum number of files per stream. | 10 | No |
| File size | The maximum file size (in MB). | 24 MB | No |
| Maximum data block size | The maximum data block size. | 128 KB | No |
| Minimum data block size | The minimum data block size. | 256 bytes | No |
| Maximum block offset specified in a stream file request | The maximum block offset specified in a stream file request. | 98,304 | No |
| Maximum blocks that can be requested per stream file request | The maximum number of blocks that can be requested per stream file request. | 98,304 | No |
| Maximum block bitmap size | The maximum block bitmap size. | 12,288 bytes | No |

* For additional information, see Using AWS IoT MQTT-based file delivery in devices in the *AWS IoT Developer Guide*.

## MQTT-based File Delivery Throttling

| API | Transactions Per Second |
|---|---|
| CreateStream | 15 TPS |

| API | Transactions Per Second |
| --- | --- |
| DeleteStream | 15 TPS |
| DescribeStream | 15 TPS |
| ListStreams | 15 TPS |
| UpdateStream | 15 TPS |
| GetStream (MQTT-based) | 1000 TPS |
| DescribeStream (MQTT-based) | 1000 TPS |

## AWS IoT Core Device Advisor limits and quotas

### AWS IoT Core Device Advisor limits and quotas

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Concurrently connected devices | The maximum number of test devices that can be concurrently connected per test suite run. | 1 | No |
| Concurrently running test suites | The maximum number of suites an AWS account can run concurrently. | 1 | No |
| Connections per test endpoint | The maximum number of connections to an account-specific test endpoint. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| MQTT CONNECT requests per account | The maximum number of MQTT Connect requests sent from a test device per second per account. | 5 | No |
| MQTT CONNECT requests per client ID | The maximum number of MQTT Connect requests sent from a test device per second per client ID. | 1 | No |
| Rate of CreateSuiteDefinition API requests | The maximum number of CreateSuiteDefinition API requests you can make per second. | 1 | No |
| Rate of DeleteSuiteDefinition API requests | The maximum number of DeleteSuiteDefinition API requests you can make per second. | 10 | No |
| Rate of GetSuiteDefinition API requests | The maximum number of GetSuiteDefinition API requests you can make per second. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Rate of GetSuiteRun API requests | The maximum number of GetSuiteRun API requests you can make per second. | 10 | No |
| Rate of GetSuiteRunReport API requests | The maximum number of GetSuiteRunReport API requests you can make per second. | 10 | No |
| Rate of ListSuite Definitions API requests | The maximum number of ListSuite Definitions API requests you can make per second. | 10 | No |
| Rate of ListSuiteRuns API requests | The maximum number of ListSuite Runs API requests you can make per second. | 10 | No |
| Rate of ListTagsForResource API requests | The maximum number of ListTagsForResource API requests you can make per second. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Rate of ListTestCases API requests | The maximum number of ListTestCases API requests you can make per second. | 10 | No |
| Rate of StartSuiteRun API requests | The maximum number of StartSuiteRun API requests you can make per second. | 1 | No |
| Rate of TagResource API requests | The maximum number of TagResource API requests you can make per second. | 10 | No |
| Rate of UntagResource API requests | The maximum number of UntagResource API requests you can make per second. | 10 | No |
| Rate of UpdateSuiteDefinition API requests | The maximum number of UpdateSuiteDefinition API requests you can make per second. | 10 | No |
| Test case execution time | The maximum time until a test case fails if not completed. | 10800 Seconds | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Test cases per test suite | The maximum number of test cases in one test suite. | 50 | No |

# AWS IoT Device Defender endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | iot.us-east-2.amazonaws.com | https | |
| | | iot-fips.us-east-2.api.aws | HTTPS | |
| | | iot-fips.us-east-2.amazonaws.com | https | |
| | | iot.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | iot.us-east-1.amazonaws.com | https | |
| | | iot-fips.us-east-1.api.aws | HTTPS | |
| | | iot-fips.us-east-1.amazonaws.com | https | |
| | | iot.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | iot.us-west-1.amazonaws.com | https |
| | | iot-fips.us-west-1.api.aws | HTTPS |
| | | iot-fips.us-west-1.amazonaws.com | https |
| | | iot.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | iot.us-west-2.amazonaws.com | https |
| | | iot-fips.us-west-2.api.aws | HTTPS |
| | | iot-fips.us-west-2.amazonaws.com | https |
| | | iot.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | iot.ap-east-1.amazonaws.com | HTTPS |
| | | iot.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | iot.ap-southeast-5.amazonaws.com | HTTPS |
| | | iot.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | iot.ap-south-1.amazonaws.com | HTTPS |
| | | iot.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | iot.ap-northeast-2.amazonaws.com | HTTPS |
| | | iot.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | iot.ap-southeast-1.amazonaws.com | HTTPS |
| | | iot.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | iot.ap-southeast-2.amazonaws.com | HTTPS |
| | | iot.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | iot.ap-northeast-1.amazonaws.com | HTTPS |
| | | iot.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | iot.ca-central-1.amazonaws.com | https |
| | | iot-fips.ca-central-1.api.aws | HTTPS |
| | | iot-fips.ca-central-1.amazonaws.com | https |
| | | iot.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | iot.eu-central-1.amazonaws.com | HTTPS |
| | | iot.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | iot.eu-west-1.amazonaws.com | HTTPS |
| | | iot.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | iot.eu-west-2.amazonaws.com | HTTPS |
| | | iot.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | iot.eu-west-3.amazonaws.com | HTTPS |
| | | iot.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | iot.eu-south-2.amazonaws.com | HTTPS |
| | | iot.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | iot.eu-north-1.amazonaws.com | HTTPS |
| | | iot.eu-north-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | iot.me-south-1.amazonaws.com | HTTPS |
| | | iot.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | iot.me-central-1.amazonaws.com | HTTPS |
| | | iot.me-central-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | iot.us-gov-east-1.amazonaws.com | https |
| | | iot-fips.us-gov-east-1.api.aws | HTTPS |
| | | iot-fips.us-gov-east-1.amazonaws.com | https |
| | | iot.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | iot.us-gov-west-1.amazonaws.com | https |
| | | iot-fips.us-gov-west-1.api.aws | HTTPS |
| | | iot-fips.us-gov-west-1.amazonaws.com | https |
| | | iot.us-gov-west-1.api.aws | HTTPS |

# Service quotas

### AWS IoT Device Defender audits limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Scheduled audits | The maximum number of scheduled audits. | 5 | No |
| Simultaneous in progress on-demand audits | The maximum number of simultaneous in progress on-demand audits. | 10 | No |
| Storage duration for audit findings | The maximum time, in days, that audit findings are stored after being reported. | 90 | No |

The following service quotas apply to mitigation actions and audit mitigation action tasks:

### AWS IoT Device Defender mitigation limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Mitigation actions | The maximum number of mitigation actions. | 100 | No |

### Audit mitigation action limits

| Resource | Limit |
|---|---|
| Number of audit mitigation action tasks running at the same time | 10 tasks |

| Resource | Limit |
| --- | --- |
| Retention period for audit mitigation action tasks | 90 days |

### AWS IoT Device Defender detect limits and quotas

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Behavior metric value elements for each security profile | The maximum number of behavior metric value elements (counts, IP addresses, ports) for each security profile. | 1000 | No |
| Behaviors for each security profile | The maximum number of behaviors for each security profile | 100 | No |
| Custom metrics | The maximum number of detect custom metrics. | 100 | Yes |
| Device metric minimum delay | The minimum time, in seconds, that a device must wait between sending metric reports. | 300 Seconds | Yes |
| Device metric peak reporting rate for an account | The maximum number of device-side metric reports that can be sent, per second, from all devices in an account. | 3500 | Yes |

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Metric dimensions | The maximum number of detect metric dimensions. | 10 | No |
| Security profiles for each target | The maximum number of security profiles for each target (things or thing groups in the AWS account). | 5 | No |
| Storage duration for detect metrics | The maximum time, in days, that detect metrics are stored after being ingested. | 14 | No |
| Storage duration for detect violations | The maximum time, in days, that detect violations are stored after being generated. | 30 | No |

### ML Detect limits

| Resource | Quota | Adjustable |
|---|---|---|
| Number of Detect mitigation action tasks that can be running at the same time | 5 maximum | Yes |
| Retention period for Detect mitigation action tasks | 90 days maximum | Yes |

| Resource | Quota | Adjustable |
| --- | --- | --- |
| Retention period for models (time after which models are expired) | 30 days maximum | No |

### AWS IoT Device Defender API throttling limits

This table describes the maximum number of transactions per second (TPS) that can be made to each of these AWS IoT Device Defender API actions.

**AWS IoT Device Defender API throttling limits**

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| AttachSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the AttachSecurityProfile API. | 10 | No |
| CancelAuditMitigationActionsTask API TPS | The maximum number of transactions per second (TPS) that can be made for the CancelAuditMitigationActionsTask API. | 10 | No |
| CancelAuditTask API TPS | The maximum number of transactions per second (TPS) that can be made for the CancelAuditTask API. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| CancelDetectMitigationActionsTask API TPS | The maximum number of transactions per second (TPS) that can be made for the CancelDetectMitigationActionsTask API. | 10 | No |
| CreateAuditSuppression API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateAuditSuppression API. | 10 | No |
| CreateCustomMetric API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateCustomMetric API. | 10 | No |
| CreateMitigationAction API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateMitigationAction API. | 10 | No |
| CreateScheduledAudit API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateScheduledAudit API. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| CreateSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the CreateSecurityProfile API. | 10 | No |
| DeleteAccountAuditConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteAccountAuditConfiguration API. | 5 | No |
| DeleteAuditSuppression API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteAuditSuppression API. | 10 | No |
| DeleteCustomMetric API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteCustomMetric API. | 10 | No |
| DeleteDimension API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteDimension API. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| DeleteMitigationAction API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteMitigationAction API. | 10 | No |
| DeleteScheduledAudit API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteScheduledAudit API. | 5 | No |
| DeleteSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the DeleteSecurityProfile API. | 10 | No |
| DescribeAccountAuditConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeAccountAuditConfiguration API. | 5 | No |
| DescribeAuditFinding API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeAuditFinding API. | 25 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| DescribeAuditMitigationActionsTask API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeAuditMitigationActionsTask API. | 25 | No |
| DescribeAuditSuppression API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeAuditSuppression API. | 10 | No |
| DescribeAuditTask API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeAuditTask API. | 25 | No |
| DescribeCustomMetric API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeCustomMetric API. | 25 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| DescribeDetectMitigationActionsTask API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeDetectMitigationActionsTask API. | 10 | No |
| DescribeDimension API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeDimension API. | 10 | No |
| DescribeMitigationAction API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeMitigationAction API. | 25 | No |
| DescribeScheduledAudit API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeScheduledAudit API. | 5 | No |
| DescribeSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the DescribeSecurityProfile API. | 25 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| DetachSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the DetachSecurityProfile API. | 10 | No |
| ListActiveViolations API TPS | The maximum number of transactions per second (TPS) that can be made for the ListActiveViolations API. | 10 | No |
| ListAuditFindings API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAuditFindings API. | 10 | No |
| ListAuditMitigationActionsExecutions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAuditMitigationActionsExecutions API. | 10 | No |
| ListAuditMitigationActionsTasks API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAuditMitigationActionsTasks API. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| ListAuditSuppressions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAuditSuppressions API. | 10 | No |
| ListAuditTasks API TPS | The maximum number of transactions per second (TPS) that can be made for the ListAuditTasks API. | 10 | No |
| ListCustomMetrics API TPS | The maximum number of transactions per second (TPS) that can be made for the ListCustomMetrics API. | 10 | No |
| ListDetectMitigationActionsExecutions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListDetectMitigationActionsExecutions API. | 10 | No |
| ListDetectMitigationActionsTasks API TPS | The maximum number of transactions per second (TPS) that can be made for the ListDetectMitigationActionsTasks API. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| ListDimensions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListDimensions API. | 10 | No |
| ListMetricValues API TPS | The maximum number of transactions per second (TPS) that can be made for the ListMetricValues API. | 15 | Yes |
| ListMitigationActions API TPS | The maximum number of transactions per second (TPS) that can be made for the ListMitigationActions API. | 10 | No |
| ListScheduledAudits API TPS | The maximum number of transactions per second (TPS) that can be made for the ListScheduledAudits API. | 5 | No |
| ListSecurityProfiles API TPS | The maximum number of transactions per second (TPS) that can be made for the ListSecurityProfiles API. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| ListSecurityProfilesForTarget API TPS | The maximum number of transactions per second (TPS) that can be made for the ListSecurityProfilesForTarget API. | 10 | No |
| ListTargetsForSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the ListTargetsForSecurityProfile API. | 10 | No |
| ListViolationEvents API TPS | The maximum number of transactions per second (TPS) that can be made for the ListViolationEvents API. | 10 | No |
| PutVerificationStateOnViolation API TPS | The maximum number of transactions per second (TPS) that can be made for the PutVerificationStateOnViolation API. | 10 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| StartAuditMitigationActionsTask API TPS | The maximum number of transactions per second (TPS) that can be made for the StartAuditMitigationActionsTask API. | 10 | No |
| StartDetectMitigationActionsTask API TPS | The maximum number of transactions per second (TPS) that can be made for the StartDetectMitigationActionsTask API. | 10 | No |
| StartOnDemandAuditTask API TPS | The maximum number of transactions per second (TPS) that can be made for the StartOnDemandAuditTask API. | 10 | No |
| UpdateAccountAuditConfiguration API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateAccountAuditConfiguration API. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| UpdateAuditSuppression API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateAuditSuppression API. | 10 | No |
| UpdateCustomMetric API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateCustomMetric API. | 10 | No |
| UpdateDimension API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateDimension API. | 10 | No |
| UpdateMitigationAction API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateMitigationAction API. | 10 | No |
| UpdateScheduledAudit API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateScheduledAudit API. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| UpdateSecurityProfile API TPS | The maximum number of transactions per second (TPS) that can be made for the UpdateSecurityProfile API. | 10 | No |
| ValidateSecurityProfileBehaviors API TPS | The maximum number of transactions per second (TPS) that can be made for the ValidateSecurityProfileBehaviors API. | 10 | No |

# AWS IoT Device Management endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Topics

- AWS IoT Device Management jobs and commands data endpoints
- AWS IoT Device Management - secure tunneling endpoints
- AWS IoT Device Management - Software Package Catalog endpoints
- AWS IoT Device Management - managed integrations endpoints

- AWS IoT FIPS endpoints

### AWS IoT Device Management jobs and commands data endpoints

> **Note**
>
> If you are using data plane endpoints for the first time, use the `iot:Data-ATS` endpoints. We recommend using the AWS IoT Core data plane ATS endpoint for accessing Jobs functionality, as it enables advanced features like IPv6 and AWS PrivateLink that are not available through the `iot:Jobs` endpoint.

The commands and jobs features in AWS IoT Device Management share the same data plane endpoints. These data endpoints are specific to each AWS account and AWS Region. To find the data plane endpoint for your AWS account and AWS Region, use the **describe-endpoint** CLI command shown here, or the `DescribeEndpoint` REST API.

To retrieve the AWS IoT Core `iot:Data-ATS` dual-stack endpoint, use the following command:

```
aws iot describe-endpoint --endpoint-type iot:Data-ATS
```

To retrieve the `iot:Jobs` IPv4 only endpoint, use the following command:

```
aws iot describe-endpoint --endpoint-type iot:Jobs
```

For information about the actions supported by the AWS IoT Device Management - jobs data endpoints, see AWS IoT jobs data plane operations in the *AWS IoT API Reference*.

> **Note**
>
> For the AWS IoT Device Management commands feature, the `StartCommandExecution` API is supported in the data plane.

The following table contains AWS Region-specific endpoints that AWS IoT Core supports for the commands and jobs data operations. In the **Endpoint** column, the *`account-specific-prefix`* from your account-specific endpoint replaces the *`data`* shown in the generic endpoint representation.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | data.jobs.iot.us-east-2.amazonaws.com<br>data.jobs.iot-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | data.jobs.iot.us-east-1.amazonaws.com<br>data.jobs.iot-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | data.jobs.iot.us-west-1.amazonaws.com<br>data.jobs.iot-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | data.jobs.iot.us-west-2.amazonaws.com<br>data.jobs.iot-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | data.jobs.iot.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | data.jobs.iot.ap-southeast-5.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | data.jobs.iot.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | data.jobs.iot.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | data.jobs.iot.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | data.jobs.iot.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | data.jobs.iot.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | data.jobs.iot.ca-central-1.amazonaws.com<br>data.jobs.iot-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | data.jobs.iot.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | data.jobs.iot.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | data.jobs.iot.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | data.jobs.iot.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | data.jobs.iot.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | data.jobs.iot.eu-north-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | data.jobs.iot.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | data.jobs.iot.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | data.jobs.iot.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | data.jobs.iot.us-gov-east-1.amazonaws.com<br>data.jobs.iot-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | data.jobs.iot.us-gov-west-1.amazonaws.com<br>data.jobs.iot-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## AWS IoT Device Management - secure tunneling endpoints

AWS IoT supports additional endpoints for secure tunneling.

### Secure Tunneling Management APIs Endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | api.tunneling.iot.us-east-2.amazonaws.com<br>api.iot-tunneling-fips.us-east-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | api.tunneling.iot-fips.us-east-2.amazonaws.com<br>api.iot-tunneling.us-east-2.api.aws | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | api.tunneling.iot.us-east-1.amazonaws.com<br>api.iot-tunneling-fips.us-east-1.api.aws<br>api.tunneling.iot-fips.us-east-1.amazonaws.com<br>api.iot-tunneling.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | api.tunneling.iot.us-west-1.amazonaws.com<br>api.iot-tunneling-fips.us-west-1.api.aws<br>api.tunneling.iot-fips.us-west-1.amazonaws.com<br>api.iot-tunneling.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | api.tunneling.iot.us-west-2.amazonaws.com<br>api.iot-tunneling-fips.us-west-2.api.aws<br>api.tunneling.iot-fips.us-west-2.amazonaws.com<br>api.iot-tunneling.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | api.tunneling.iot.ap-east-1.amazonaws.com<br>api.iot-tunneling.ap-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | api.tunneling.iot.ap-southeast-5.amazonaws.com<br>api.iot-tunneling.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | api.tunneling.iot.ap-south-1.amazonaws.com<br>api.iot-tunneling.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | api.tunneling.iot.ap-northeast-2.amazonaws.com<br>api.iot-tunneling.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | api.tunneling.iot.ap-southeast-1.amazonaws.com<br>api.iot-tunneling.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | api.tunneling.iot.ap-southeast-2.amazonaws.com<br>api.iot-tunneling.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.tunneling.iot.ap-northeast-1.amazonaws.com<br>api.iot-tunneling.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | api.tunneling.iot.ca-central-1.amazonaws.com<br>api.iot-tunneling-fips.ca-central-1.api.aws<br>api.tunneling.iot-fips.ca-central-1.amazonaws.com<br>api.iot-tunneling.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | api.tunneling.iot.eu-central-1.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | api.tunneling.iot.eu-west-1.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | api.tunneling.iot.eu-west-2.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | api.tunneling.iot.eu-south-1.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | api.tunneling.iot.eu-west-3.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | api.tunneling.iot.eu-south-2.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | api.tunneling.iot.eu-north-1.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.eu-north-1.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | api.tunneling.iot.il-central-1.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.il-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | api.tunneling.iot.me-south-1.amazonaws.com | HTTPS | |
| | | api.iot-tunneling.me-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | api.tunneling.iot.me-central-1.amazonaws.com<br>api.iot-tunneling.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | api.tunneling.iot.sa-east-1.amazonaws.com<br>api.iot-tunneling.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | api.tunneling.iot.us-gov-east-1.amazonaws.com<br>api.iot-tunneling-fips.us-gov-east-1.api.aws<br>api.tunneling.iot-fips.us-gov-east-1.amazonaws.com<br>api.iot-tunneling.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | api.tunneling.iot.us-gov-west-1.amazonaws.com<br>api.iot-tunneling-fips.us-gov-west-1.api.aws<br>api.tunneling.iot-fips.us-gov-west-1.amazonaws.com<br>api.iot-tunneling.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

### Secure Tunneling Device Connection Endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | wss://data.tunneling.iot.us-east-2.amazonaws.com<br>wss://data.iot-tunneling.us-east-2.api.aws<br>wss://data.tunneling.iot-fips.us-east-2.amazonaws.com<br>wss://data.iot-tunneling-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | wss://data.tunneling.iot.us-east-1.amazonaws.com<br>wss://data.iot-tunneling.us-east-1.api.aws<br>wss://data.tunneling.iot-fips.us-east-1.amazonaws.com<br>wss://data.iot-tunneling-fips.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | wss://data.tunneling.iot.us-west-1.amazonaws.com<br>wss://data.iot-tunneling.us-west-1.api.aws<br>wss://data.tunneling.iot-fips.us-west-1.amazonaws.com<br>wss://data.iot-tunneling-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | wss://data.tunneling.iot.us-west-2.amazonaws.com<br>wss://data.iot-tunneling.us-west-2.api.aws<br>wss://data.tunneling.iot-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | wss://data.iot-tunneling-fips.us-west-2.api.aws | |
| Asia Pacific (Malaysia) | ap-southeast-5 | wss://data.tunneling.iot.ap-southeast-5.amazonaws.com<br>wss://data.iot-tunneling.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | wss://data.tunneling.iot.ap-south-1.amazonaws.com<br>wss://data.iot-tunneling.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | wss://data.tunneling.iot.ap-northeast-2.amazonaws.com<br>wss://data.iot-tunneling.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | wss://data.tunneling.iot.ap-southeast-1.amazonaws.com<br>wss://data.iot-tunneling.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | wss://data.tunneling.iot.ap-southeast-2.amazonaws.com<br>wss://data.iot-tunneling.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | wss://data.tunneling.iot.ap-northeast-1.amazonaws.com<br>wss://data.iot-tunneling.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | wss://data.tunneling.iot.ap-east-1.amazonaws.com<br>wss://data.iot-tunneling.ap-east-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | wss://data.tunneling.iot.ca-central-1.amazonaws.com<br>wss://data.iot-tunneling.ca-central-1.api.aws<br>wss://data.tunneling.iot-fips.ca-central.1.amazonaws.com<br>wss://data.iot-tunneling-fips.ca-central.1.api.aws | HTTPS |
| China (Beijing) | cn-north-1 | wss://data.tunneling.iot.cn-north-1.amazonaws.com.cn<br>wss://data.iot-tunneling.cn-north-1.api.amazonwebservices.com.cn | HTTPS |
| China (Ningxia) | cn-northwest-1 | wss://data.tunneling.iot.cn-northwest-1.amazonaws.com.cn<br>wss://data.iot-tunneling.cn-northwest-1.api.amazonwebservices.com.cn | HTTPS |
| Europe (Frankfurt) | eu-central-1 | wss://data.tunneling.iot.eu-central-1.amazonaws.com<br>wss://data.iot-tunneling.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | wss://data.tunneling.iot.eu-west-1.amazonaws.com<br>wss://data.iot-tunneling.eu-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | wss://data.tunneling.iot.eu-west-2.amazonaws.com<br>wss://data.iot-tunneling.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | wss://data.tunneling.iot.eu-west-3.amazonaws.com<br>wss://data.iot-tunneling.eu-west-3.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | wss://data.tunneling.iot.eu-south-1.amazonaws.com<br>wss://data.iot-tunneling.eu-south-1.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | wss://data.tunneling.iot.eu-south-2.amazonaws.com<br>wss://data.iot-tunneling.eu-south-2.api.aws | |
| Europe (Stockholm) | eu-north-1 | wss://data.tunneling.iot.eu-north-1.amazonaws.com<br>wss://data.iot-tunneling.eu-north-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | wss://data.tunneling.iot.sa-east-1.amazonaws.com<br>wss://data.iot-tunneling.sa-east-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | wss://data.tunneling.iot.me-south-1.amazonaws.com<br>wss://data.iot-tunneling.me-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | wss://data.tunneling.iot.me-central-1.amazonaws.com<br><br>wss://data.iot-tunneling.me-central-1.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | wss://data.tunneling.iot.il-central-1.amazonaws.com<br><br>wss://data.iot-tunneling.il-central-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | wss://data.iot-tunneling.us-gov-east-1.api.aws<br><br>wss://data.tunneling.iot-fips.us-gov-east-1.amazonaws.com<br><br>wss://data.iot-tunneling-fips.us-gov-east-1.api.aws<br><br>wss://data.tunneling.iot-fips.us-gov-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | wss://data.iot-tunneling.us-gov-west-1.api.aws<br><br>wss://data.tunneling.iot-fips.us-gov-west-1.amazonaws.com<br><br>wss://data.iot-tunneling-fips.us-gov-west-1.api.aws<br><br>wss://data.tunneling.iot-fips.us-gov-west-1.amazonaws.com | HTTPS |

## AWS IoT Device Management - Software Package Catalog endpoints

AWS IoT Device Management Software Package Catalog is available in the following AWS Regions.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | iot.us-east-2.amazonaws.com<br>iot-fips.us-east-2.amazonaws.com<br>iot.us-east-2.api.aws<br>iot-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | iot.us-east-1.amazonaws.com<br>iot-fips.us-east-1.amazonaws.com<br>iot.us-east-1.api.aws<br>iot-fips.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | iot.us-west-1.amazonaws.com<br>iot-fips.us-west-1.amazonaws.com<br>iot.us-west-1.api.aws<br>iot-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | iot.us-west-2.amazonaws.com<br>iot-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | iot.us-west-2.api.aws<br><br>iot-fips.us-west-2.api.aws | |
| Asia Pacific (Hong Kong) | ap-east-1 | iot.ap-east-1.amazonaws.com<br><br>iot.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | iot.ap-southeast-5.amazonaws.com<br><br>iot.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | iot.ap-south-1.amazonaws.com<br><br>iot.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | iot.ap-northeast-2.amazonaws.com<br><br>iot.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | iot.ap-southeast-1.amazonaws.com<br><br>iot.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | iot.ap-southeast-2.amazonaws.com<br><br>iot.ap-southeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | iot.ap-northeast-1.amazonaws.com<br>iot.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | iot.ca-central-1.amazonaws.com<br>iot-fips.ca-central-1.amazonaws.com<br>iot.ca-central-1.api.aws<br>iot-fips.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | iot.eu-central-1.amazonaws.com<br>iot.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | iot.eu-west-1.amazonaws.com<br>iot.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | iot.eu-west-2.amazonaws.com<br>iot.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | iot.eu-west-3.amazonaws.com<br>iot.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Milan) | eu-south-1 | iot.eu-south-1.amazonaws.com<br>iot.eu-south-1.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | iot.eu-south-2.amazonaws.com<br>iot.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | iot.eu-north-1.amazonaws.com<br>iot.eu-north-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | iot.me-south-1.amazonaws.com<br>iot.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | iot.me-central-1.amazonaws.com<br>iot.me-central-1.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | iot.il-central-1.amazonaws.com<br>iot.il-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| South America (São Paulo) | sa-east-1 | iot.sa-east-1.amazonaws.com<br>iot.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | iot.us-gov-east-1.amazonaws.com<br>iot-fips.us-gov-east-1.amazonaws.com<br>iot.us-gov-east-1.api.aws<br>iot-fips.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | iot.us-gov-west-1.amazonaws.com<br>iot-fips.us-gov-west-1.amazonaws.com<br>iot.us-gov-west-1.api.aws<br>iot-fips.us-gov-west-1.api.aws | HTTPS |
| China (Beijing) | cn-north-1 | iot.cn-north-1.amazonaws.com.cn<br>iot.cn-north-1.api.amazonwebservices.com.cn | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| China (Ningxia) | cn-northwest-1 | iot.cn-northwest-1.amazonaws.com.cn<br><br>iot.cn-northwest-1.api.amazonwebservices.com.cn | HTTPS |

## AWS IoT Device Management - managed integrations endpoints

Managed integrations for AWS IoT Device Management is available in the following AWS Regions.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | api.iotmanagedintegrations.ca-central-1.api.aws<br><br>api.iotmanagedintegrations-fips.ca-central-1.api.aws | HTTPS<br><br>HTTPS |
| Europe (Ireland) | eu-west-1 | api.iotmanagedintegrations.eu-west-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | api.iotmanagedintegrations.me-central-1.api.aws | HTTPS |

## AWS IoT FIPS endpoints

AWS IoT provides endpoints that support the Federal Information Processing Standard (FIPS) 140-2. Choose the appropriate FIPS compliant endpoint to access AWS IoT features in your AWS Region from FIPS Endpoints by Service. For more information about the FIPS endpoints provided by AWS IoT, see Connecting to AWS IoT FIPS endpoints.

# Service quotas

### Contents


- AWS IoT Fleet Indexing
- AWS IoT Jobs
- AWS IoT Device Management commands
- AWS IoT Secure Tunneling
- AWS IoT Device Management - Software Package Catalog
- AWS IoT Device Management - managed integrations


## AWS IoT Fleet Indexing

### AWS IoT Device Management fleet indexing limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Maximum length of a custom field name | The maximum length of a custom field name. | 1024 | Yes |
| Maximum length of a query | The maximum length of a query in UTF-8 encoded characters. | 1000 | Yes |
| Maximum number of * wildcard operators per query term | The maximum number of * wildcard operators per query term. | 2 | No |
| Maximum number of ? wildcard | The maximum number of ? wildcard operators per query term. | 5 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| operators per query term | | | |
| Maximum number of custom fields in AWS thing groups index | The maximum number of custom fields in AWS thing groups index. | 5 | Yes |
| Maximum number of custom fields in AWS things index | The maximum number of custom fields in AWS things index. | 5 | Yes |
| Maximum number of dynamic groups | The maximum number of dynamic groups per customer. | 100 | No |
| Maximum number of fleet metrics | The maximum number of fleet metrics per customer. | 100 | Yes |
| Maximum number of names in the named shadow names filter | The maximum number of names in the named shadow names filter | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum number of percentile values per fleet metric | The maximum number of values for percentile aggregation type per fleet metric. | 5 | No |
| Maximum number of query terms per dynamic group | The maximum number of query terms per dynamic group. | 5 | Yes |
| Maximum number of query terms per fleet metric | The maximum number of query terms per fleet metric. | 12 | Yes |
| Maximum number of query terms per query | The maximum number of query terms per query. | 12 | Yes |
| Maximum number of results per search query results page | The maximum number of results displayed per search query results page. | 500 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum number of targets in the geo locations filter | The maximum number of targets in the geo locations filter | 1 | Yes |
| Maximum period of a fleet metric | The maximum period of a fleet metric in seconds. | 86400 Seconds | No |
| Minimum period of a fleet metric | The minimum period of a fleet metric in seconds. | 60 Seconds | No |

### AWS IoT Device Management fleet indexing API limits

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| DescribeIndex rate | The maximum number of DescribeI ndex calls per second. | 10 | Yes |
| GetCardinality rate | The maximum number of GetCardin ality calls per second. | 15 | Yes |
|  | The maximum number of GetIndexi ngConfiguration calls per second. | 20 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| GetIndexingConfiguration rate | | | |
| GetPercentiles rate | The maximum number of GetPercentiles calls per second. | 15 | Yes |
| GetStatistics rate | The maximum number of GetStatistics calls per second. | 15 | Yes |
| ListIndices rate | The maximum number of ListIndices calls per second. | 5 | Yes |
| SearchIndex rate | The maximum number of SearchIndex calls per second. | 15 | Yes |
| GetThingConnectivityData rate | The maximum number of GetThingConnectivityData calls per second. | 350 | Yes |
| UpdateIndexingConfiguration rate | The maximum number of UpdateIndexingConfiguration calls per second. | 1 | Yes |

# AWS IoT Jobs

## AWS IoT Device Management jobs limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| Active continuous jobs | The maximum number of active Jobs in your AWS account per region. This limit value is an aggregation of both active snapshot jobs and active continuous jobs. | 100000 | Yes |
| Active snapshot jobs | The maximum number of active Jobs in your AWS account per region. This limit value is an aggregation of both active snapshot jobs and active continuous jobs. | 100000 | Yes |
| Comment length | The maximum comment length (in characters). | 2028 | No |
| Concurrent jobs | The maximum number of concurrent jobs in your AWS account per region. | 500 | Yes |
| Data retention | The maximum number of days that job data and job execution data will be | 730 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| | retained for inactive jobs (jobs that aren't IN_PROGRESS). | | |
| DocumentSource length | The maximum number of characters in a job document source. | 1350 | No |
| In Progress timeout | The maximum job execution InProgress timeout value (in minutes). | 10080 | No |
| Job Targets | The maximum number of targets you can assign to a job. | 100 | No |
| Job Template description length | The maximum number of characters in a job template description. | 2028 | No |
| Job description length | The maximum number of characters in a job description. | 2028 | No |
| Job execution roll out rate | The maximum number of job executions that you can roll out per minute. | 1000 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| JobId Length | The maximum number of characters in a Job id. | 64 | No |
| JobTemplateId Length | The maximum number of characters in a job template id. | 64 | No |
| List results per page | The maximum number of list results per page. | 250 | No |
| Maximum number of job templates | The maximum number of job templates you can own. | 100 | Yes |
| Minimum job execution roll out rate | The minimum number of job executions that you can roll out per minute. | 1 | No |
| Minimum pre-signed URL lifetime | The minimum lifetime (in seconds) of a pre-signed URL. | 60 Seconds | No |
| Pre-signed URL lifetime | The maximum lifetime (in seconds) of a pre-signed URL. | 3600 Seconds | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| S3 job document length | The maximum length of an S3 job document that can be sent to an AWS IoT device (in Bytes). | 32768 Bytes | Yes |
| StatusDetail map key length | The maximum length of a StatusDetail map key (in characters). | 128 | No |
| StatusDetail map key-value pairs | The maximum number of key-value pairs you can have in a StatusDetail map. | 10 | No |
| StatusDetail map value length | The maximum length of a StatusDetail map value (in characters). | 1024 | No |
| Step Timer | The maximum job execution step timeout value (in minutes). | 10080 | No |

### AWS IoT Device Management jobs API action limits

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| AssociateTargetsWithJob throttle limit | The throttle limit for AssociateTargetsWithJob in your AWS account per Region per second. | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| CancelJob throttle limit | The throttle limit for CancelJob in your AWS account per Region per second. | 10 | Yes |
| CancelJob Execution throttle limit | The throttle limit for CancelJobExecution in your AWS account per Region per second. | 10 | Yes |
| CreateJob throttle limit | The throttle limit for CreateJob in your AWS account per Region per second. | 10 | No |
| CreateJob Template throttle limit | The throttle limit for CreateJobTemplate in your AWS account per Region per second. | 10 | Yes |
| DeleteJob throttle limit | The throttle limit for DeleteJob in your AWS account per Region per second. | 10 | Yes |
| DeleteJob Execution throttle limit | The throttle limit for DeleteJobExecution in your AWS account per Region per second. | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| DeleteJobTemplate throttle limit | The throttle limit for DeleteJobTemplate in your AWS account per Region per second. | 10 | Yes |
| DescribeJob throttle limit | The throttle limit for DescribeJob in your AWS account per Region per second. | 10 | Yes |
| DescribeJobExecution throttle limit | The throttle limit for DescribeJobExecution in your AWS account per Region per second. | 10 | Yes |
| DescribeJobExecution/GetPendingJobExecutions throttle limit | The maximum number of total read transactions per second per account which can be caused by invoking DescribeJobExecution and/or GetPendingJobExecutions. In the control plane, DescribeJobExecution is limited to 10 TPS per invocation. | 200 | No |
| DescribeJobTemplate throttle limit | The throttle limit for DescribeJobTemplate in your AWS account per Region per second. | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| GetJobDocument throttle limit | The throttle limit for GetJobDocument in your AWS account per Region per second. | 10 | Yes |
| ListJobExecutionsForJob throttle limit | The throttle limit for ListJobExecutionsForJob in your AWS account per Region per second. | 10 | Yes |
| ListJobExecutionsForThing throttle limit | The throttle limit for ListJobExecutionsForThing in your AWS account per Region per second. | 10 | Yes |
| ListJobTemplates throttle limit | The throttle limit for ListJobTemplates in your AWS account per Region per second. | 10 | Yes |
| ListJobs throttle limit | The throttle limit for ListJobs in your AWS account per Region per second. | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| StartNextPendingJobExecution/UpdateJobExecution throttle limit | The maximum number of total write transactions per second per account which can be caused by invoking StartNextPendingJobExecution and/or UpdatePendingJobExecution. | 200 | No |
| UpdateJob throttle limit | The throttle limit for UpdateJob in your AWS account per Region per second. | 10 | Yes |

† For definitions of data plane and control plane, see **What are the ways for accessing AWS IoT Core?** in the AWS IoT Core FAQs

## AWS IoT Device Management commands

### AWS IoT Device Management commands limits and quotas

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Command execution timeout | The command execution timeout in seconds in your AWS account | 43200 | Yes |
| Retention period for command executions | The number of days for which the command executions must be retained in your account. | 180 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Command execution concurrency limit | The maximum number of concurrent command executions in your AWS account for each device. | 100 | No |

### AWS IoT Device Management commands API action limits

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| CreateCommand throttle limit | The throttle limit for CreateCommand in your AWS account per Region per second. | 10 | Yes |
| DeleteCommand throttle limit | The throttle limit for DeleteCommand in your AWS account per Region per second. | 10 | Yes |
| DeleteCommandExecution throttle limit | The throttle limit for DeleteCommandExecution in your AWS account per Region per second. | 15 | Yes |
| GetCommand throttle limit | The throttle limit for GetCommand in your AWS account per Region per second. | 15 | Yes |
| | The throttle limit for GetCommandExecution in your | 20 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| GetCommandExecution throttle limit | AWS account per Region per second. | | |
| ListCommands throttle limit | The throttle limit for ListCommands in your AWS account per Region per second. | 10 | Yes |
| ListCommandExecutions throttle limit | The throttle limit for ListCommandExecutions in your AWS account per Region per second. | 10 | Yes |
| StartCommandExecution throttle limit | The throttle limit for StartCommandExecution in your AWS account per Region per second. | 100 | Yes |
| UpdateCommand throttle limit | The throttle limit for UpdateCommand in your AWS account per Region per second. | 10 | Yes |
| UpdateCommandExecution throttle limit | The throttle limit for UpdateCommandExecution in your AWS account per Region per second. | 100 | Yes |

## AWS IoT Secure Tunneling

### AWS IoT Device Management secure tunneling limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| CloseTunnel API throttle limit | The maximum number of transactions per second per account which can be caused by invoking CloseTunnel. | 1 | Yes |
| DescribeTunnel API throttle limit | The maximum number of transactions per second per account which can be caused by invoking DescribeTunnel. | 10 | Yes |
| ListTagsForResource API throttle limit | The maximum number of transactions per second per account which can be caused by invoking ListTagsForResource. | 10 | Yes |
| ListTunnels API throttle limit | The maximum number of transactions per second per account which can be caused by invoking ListTunnels. | 10 | Yes |
| Maximum bandwidth per tunnel | The maximum bandwidth per tunnel (in kbps). | 800 | No |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| Maximum connection rate | The maximum number of transactions for connecting to a tunnel per second. | 10 | Yes |
| Maximum number of tags per resource | The maximum number of tags that can be used per resource. | 50 | No |
| Maximum tag key length | The maximum number of Unicode characters in a tag key. Each resource and tag key must be unique. | 128 | No |
| Maximum tag value length | The maximum number of Unicode characters in a tag value. Each tag key can have one value. | 256 | No |
| Maximum tunnel lifetime | The maximum tunnel lifetime (in hours), after which a tunnel will be closed after reaching. | 12 | No |
| OpenTunnel API throttle limit | The maximum number of transactions per second per account which can be caused by invoking OpenTunnel. | 1 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TagResource API throttle limit | The maximum number of transactions per second per account which can be caused by invoking TagResource. | 10 | Yes |
| UntagResource API trottle limit | The maximum number of transactions per second per account which can be caused by invoking UntagResource. | 10 | Yes |

## AWS IoT Device Management - Software Package Catalog

### Software Package Catalog software package and package version limits

| Name | Description | Default | Adjustable |
| --- | --- | --- | --- |
| Packages per account | The maximum number of packages per account. | 100 | Yes |
| Package versions per account | The maximum number of package versions per account. | 10000 | Yes |
| Package name length | The maximum package name length. | 128 bytes | No |
| Package version name length | The maximum package version name length. | 64 bytes | No |

| Name | Description | Default | Adjustable |
| --- | --- | --- | --- |
| Package description length | The maximum length of the description field for packages. | 1024 bytes | No |
| Package version description length | The maximum length of the description field for package versions. | 1024 bytes | No |
| Size of package version attributes | The maximum size of attributes that capture additiona l metadata for a package version. | 3K bytes | No |

### Software Package Catalog API action limits

| Name | Description | Default | Adjustable |
| --- | --- | --- | --- |
| CreatePackage API TPS | The maximum CreatePackage API transactions per second. | 10 | Yes |
| CreatePackageVersi on API TPS | The maximum CreatePackageVersi on API transactions per second. | 10 | Yes |
| DeletePackage API TPS | The maximum DeletePackage API transactions per second. | 10 | Yes |
| DeletePackageVersi on API TPS | The maximum DeletePackageVersi | 10 | Yes |

| Name | Description | Default | Adjustable |
| --- | --- | --- | --- |
| | on API transactions per second. | | |
| GetPackage API TPS | The maximum GetPackage API transactions per second. | 10 | Yes |
| GetPackageVersion API TPS | The maximum GetPackageVersion API transactions per second. | 10 | Yes |
| ListPackages API TPS | The maximum ListPackages API transactions per second. | 10 | Yes |
| UpdatePackage API TPS | The maximum UpdatePackage API transactions per second. | 10 | Yes |
| UpdatePackageVersion API TPS | The maximum UpdatePackageVersion API transactions per second. | 10 | Yes |
| GetPackageConfiguration API TPS | The maximum GetPackageConfiguration API transactions per second. | 5 | Yes |
| UpdatePackageConfiguration API TPS | The maximum UpdatePackageConfiguration API transactions per second. | 5 | Yes |

## AWS IoT Device Management - managed integrations

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateAccountAssociation requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateAccountAssociation API. |
| Rate of CreateCloudConnector requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateCloudConnector API. |
| Rate of CreateConnectorDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateConnectorDestination API. |
| Rate of CreateCredentialLocker requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for CreateCredentialLocker API. |
| Rate of CreateDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateDestination API. |
| Rate of CreateEventLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | made for the CreateEventLogConfiguration API. |
| Rate of CreateManagedThing requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateManagedThing API. |
| Rate of CreateNotificationConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateNotificationConfiguration API. |
| Rate of CreateOtaTask requests | Each supported Region: 5 | Yes | The maximum number of transactions per second (TPS) that can be made for CreateOtaTask API. |
| Rate of CreateOtaTaskConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateOtaTaskConfiguration API. |
| Rate of CreateProvisioningProfile requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the CreateProvisioningProfile API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteAccountAssociation requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteAccountAssociation API. |
| Rate of DeleteCloudConnector requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteCloudConnector API. |
| Rate of DeleteConnectorDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteConnectorDestination API. |
| Rate of DeleteCredentialLocker requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for DeleteCredentialLocker API. |
| Rate of DeleteDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteDestination API. |
| Rate of DeleteEventLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteEventLogConfiguration API. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DeleteManagedThing requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteManagedThing API. |
| Rate of DeleteNotificationConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteNotificationConfiguration API. |
| Rate of DeleteOtaTask requests | Each supported Region: 5 | Yes | The maximum number of transactions per second (TPS) that can be made for DeleteOtaTask API. |
| Rate of DeleteOtaTaskConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteOtaTaskConfiguration API. |
| Rate of DeleteProvisioningProfile requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeleteProvisioningProfile API. |
| Rate of DeregisterAccountAssociation requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the DeregisterAccountAssociation API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetAccountAssociation requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetAccountAssociation API. |
| Rate of GetCloudConnector requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetCloudConnector API. |
| Rate of GetConnectorDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetConnectorDestination API. |
| Rate of GetCredentialLocker requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for GetCredentialLocker API. |
| Rate of GetCustomEndpoint requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetCustomEndpoint API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetDefaultEncryptionConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetDefaultEncryptionConfiguration API. |
| Rate of GetDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetDestination API. |
| Rate of GetDeviceDiscovery requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetDeviceDiscovery API. |
| Rate of GetEventLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetEventLogConfiguration API. |
| Rate of GetHubConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetHubConfiguration API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetManagedThing requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetManagedThing API. |
| Rate of GetManagedThingCapabilities requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetManagedThingCapabilities API. |
| Rate of GetManagedThingCertificate requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetManagedThingCertificate API. |
| Rate of GetManagedThingConnectivityData requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for GetManagedThingConnectivityData API. |
| Rate of GetManagedThingMetaData requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetManagedThingMetaData API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetManagedThingState requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetManagedThingState API. |
| Rate of GetNotificationConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetNotificationConfiguration API. |
| Rate of GetOtaTask requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for GetOtaTask API. |
| Rate of GetOtaTaskConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for GetOtaTaskConfiguration API. |
| Rate of GetProvisioningProfile requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetProvisioningProfile API. |
| Rate of GetRuntimeLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the GetRuntimeLogConfiguration API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetSchemaVersion requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for GetSchemaVersion API. |
| Rate of ListAccountAssociations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListAccountAssociations API. |
| Rate of ListCloudConnectors requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListCloud Connectors API. |
| Rate of ListConnectorDestinations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListConnectorDestinations API. |
| Rate of ListCredentialLockers requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for ListCredentialLockers API. |
| Rate of ListDestinations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListDestinations API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListDeviceDiscoveries requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListDeviceDiscoveries API. |
| Rate of ListDiscoveredDevices requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListDiscoveredDevices API. |
| Rate of ListEventLogConfigurations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListEventLogConfigurations API. |
| Rate of ListManagedThingAccountAssociations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListManagedThingAccountAssociations API. |
| Rate of ListManagedThingSchemas requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for ListManagedThingSchemas API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListManagedThings requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListManagedThings API. |
| Rate of ListNotificationConfigurations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListNotificationConfigurations API. |
| Rate of ListOtaTaskConfigurations requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for ListOtaTaskConfigurations API. |
| Rate of ListOtaTaskExecutions requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for ListOtaTaskExecutions API. |
| Rate of ListOtaTasks requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for ListOtaTasks API. |
| Rate of ListProvisioningProfiles requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListProvisioningProfiles API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListSchemaVersions requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for ListSchemaVersions API. |
| Rate of ListTagsForResource requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ListTagsForResource API. |
| Rate of PutDefaultEncryptionConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the PutDefaultEncryptionConfiguration API. |
| Rate of PutHubConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the PutHubConfiguration API. |
| Rate of PutRuntimeLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the PutRuntimeLogConfiguration API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of RegisterAccountAssociation requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the RegisterAccountAssociation API. |
| Rate of RegisterCustomEndpoint requests | Each supported Region: 1 | Yes | The maximum number of transactions per second (TPS) that can be made for the RegisterCustomEndpoint API. |
| Rate of ResetRuntimeLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the ResetRuntimeLogConfiguration API. |
| Rate of SendConnectorEvent requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the SendConnectorEvent API. |
| Rate of SendManagedThingCommand requests | Each supported Region: 100 | Yes | The maximum number of transactions per second (TPS) that can be made for the SendManagedThingCommand API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartAccountAssociationRefresh requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the StartAccountAssociationRefresh API. |
| Rate of StartDeviceDiscovery requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the StartDeviceDiscovery API. |
| Rate of TagResource requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the TagResource API. |
| Rate of UntagResource requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UntagResource API. |
| Rate of UpdateAccountAssociation requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateAccountAssociation API. |
| Rate of UpdateCloudConnector requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateCloudConnector API. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateConnectorDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateConnectorDestination API. |
| Rate of UpdateDestination requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateDestination API. |
| Rate of UpdateEventLogConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateEventLogConfiguration API. |
| Rate of UpdateManagedThing requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateManagedThing API. |
| Rate of UpdateNotificationConfiguration requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for the UpdateNotificationConfiguration API. |
| Rate of UpdateOtaTask requests | Each supported Region: 10 | Yes | The maximum number of transactions per second (TPS) that can be made for UpdateOtaTask API. |

# AWS IoT FleetWise endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Important**

Access to certain AWS IoT FleetWise features is currently gated. For more information, see AWS Region and feature availability in the *AWS IoT FleetWise Developer Guide*.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | iotfleetwise.us-east-1.amazonaws.com | HTTPS |
| | | iotfleetwise.us-east-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | iotfleetwise.ap-south-1.amazonaws.com | HTTPS |
| | | iotfleetwise.ap-south-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | iotfleetwise.eu-central-1.amazonaws.com | HTTPS |
| | | iotfleetwise.eu-central-1.api.aws | HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum size of a message | Each supported Region: 128 Kilobytes | No | The maximum message size that AWS IoT FleetWise can ingest or send to the vehicle. |
| Number of campaigns for each account in an AWS Region | Each supported Region: 20 | Yes | The maximum number of campaigns for each account in an AWS Region. |
| Number of data dimensions for a state template in an AWS Region. | Each supported Region: 5 | Yes | The maximum number of data dimensions for a state template in an AWS Region. |
| Number of decoder manifests for a model manifest for each account in an AWS Region. | Each supported Region: 100 | Yes | The maximum number of decoder manifests for a model manifest for each account in an AWS Region. |
| Number of metadata dimensions for a state template in an AWS Region. | Each supported Region: 5 | Yes | The maximum number of metadata dimensions for a state template in an AWS Region. |
| Number of model manifests for each account in an AWS Region. | Each supported Region: 150 | Yes | The maximum number of model manifests for each account in an AWS Region. |
| Number of nodes in a signal catalog for each account in an AWS Region | Each supported Region: 5,000 | Yes | The maximum number of nodes in a signal catalog |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | for each account in an AWS Region. |
| Number of partitions in a campaign configured to store and forward. | Each supported Region: 5 | Yes | The maximum number of partitions that ingest data for each campaign configured to store and forward. |
| Number of signal catalogs for each account in an AWS Region. | Each supported Region: 1 | No | The maximum number of signal catalogs for each account in an AWS Region. |
| Number of signals for a state template in an AWS Region. | Each supported Region: 500 | Yes | The maximum number of signals in a state template in an AWS Region. |
| Number of signals in a campaign for each account in an AWS Region. | Each supported Region: 500 | Yes | The maximum number of signals in a campaign for each account in an AWS Region. |
| Number of state templates for each account in an AWS Region. | Each supported Region: 20 | Yes | The maximum number of state templates for each account in an AWS Region. |
| Number of state templates for each vehicle in an AWS Region. | Each supported Region: 20 | Yes | The maximum number of state templates for each vehicle in an AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of vehicles in a fleet for each account in an AWS Region. | Each supported Region: 2,000 | Yes | The maximum number of vehicles in a fleet for each account in an AWS Region. |
| Rate of API requests for each account in an AWS Region | Each supported Region: 20 | Yes | The maximum number of API requests that you can send per second for each account in an AWS Region. |
| Rate of check-ins for each vehicle in an AWS Region | Each supported Region: 6 | No | The maximum number of check-ins that a vehicle can send per minute in an AWS Region. |
| Rate of ingesting messages for each account in an AWS Region. | Each supported Region: 1,000 | Yes | The maximum number of messages AWS IoT FleetWise can ingest per second for each account in an AWS Region. |
| Rate of ingesting messages for each vehicle in an AWS Region. | Each supported Region: 1 | Yes | The maximum number of messages AWS IoT FleetWise can ingest per second for each vehicle in an AWS Region. |

The following are service limits for data plane vision system data.

**Note**

Vision system data is in preview release and is subject to change.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of data | 128 | Yes | The maximum capacity in MB per second for messages and Amazon Ion data ingested into AWS IoT FleetWise. |
| Size of Ion object | 100 | No | The maximum size in MB of a single Ion file ingested into AWS IoT FleetWise. |
| Size of a signal sample | 200 | No | The maximum size in KB of metadata (typed data) provided in a signal sample in an Ion file. |
| Size of all signal samples | 2 | No | The maximum size in MB of all metadata (typed data) provided across all sample signals in an Ion file. |
| Rate of structured data | 10 | Yes | The maximum capacity in MB per second of all metadata (typed data) ingested across all Ion files. |
| Rate of binary data | 1,000 | Yes | The maximum number of binary blobs per second |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | ingested into AWS IoT FleetWise. |

**Note**

AWS IoT FleetWise will drop messages from the vehicle if the ingest rate exceeds the default quota. Revisit your campaign and fleet definitions to adjust the expected number of messages. Any dropped messages cannot be recovered.

# AWS IoT Wireless endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The following sections describe the service endpoints for AWS IoT Wireless. AWS IoT Core for Amazon Sidewalk can only be used in the AWS Region `us-east-1`. You can use these endpoints to perform the operations in the AWS IoT Wireless API Reference.

For information about:

- Connecting to and using the LoRaWAN gateways and devices, see Connecting gateways and devices to AWS IoT Core for LoRaWAN in the *AWS IoT Developer Guide*.
- Connecting to and using the Amazon Sidewalk end devices, see Connecting Sidewalk end devices to AWS IoT Core for Amazon Sidewalk in the *AWS IoT Developer Guide*.

### Topics

- AWS IoT Core for LoRaWAN - control plane endpoints
- AWS IoT Core for LoRaWAN - data plane endpoints

## AWS IoT Core for LoRaWAN - control plane endpoints

The following table contains AWS Region-specific endpoints for AWS IoT Core for LoRaWAN - control plane operations. You can use the regular endpoint which has the format `api.iotwireless.<region>.amazonaws.com`, or the dual-stack endpoint, which uses the format `api.iotwireless.<region>.api.aws`.

**Note**

- When using the endpoints for control plane clients, you must provide the Server Name Indication (SNI) extension. The clients can use the SNI extension to indicate the name of the server being contacted, and whether it's using the regular endpoints or the dual-stack endpoints.
- For AWS IoT Core for Amazon Sidewalk, use the endpoint in the `us-east-1` AWS Region.

For information about the operations supported by the AWS IoT Wireless - control plane endpoints, see AWS IoT Wireless API operations in the *AWS IoT Wireless API Reference*.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | api.iotwireless.us-east-1.amazonaws.com<br>api.iotwireless.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | api.iotwireless.us-west-2.amazonaws.com<br>api.iotwireless.us-west-2.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | api.iotwireless.eu-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | api.iotwireless.eu-west-1.api.aws | |
| Europe (Frankfurt) | eu-central-1 | api.iotwireless.eu-central-1.amazonaws.com<br><br>api.iotwireless.eu-central-1.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.iotwireless.ap-northeast-1.amazonaws.com<br><br>api.iotwireless.ap-northeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | api.iotwireless.ap-southeast-2.amazonaws.com<br><br>api.iotwireless.ap-southeast-2.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | api.iotwireless.sa-east-1.amazonaws.com<br><br>api.iotwireless.sa-east-1.api.aws | HTTPS |

## AWS IoT Core for LoRaWAN - data plane endpoints

The AWS IoT Core for LoRaWAN - data plane endpoints are specific to each AWS account and AWS Region. To find the AWS IoT Core for LoRaWAN - data plane endpoint for your AWS account and AWS Region, use the get-service-endpoint CLI command shown here, or the GetServiceEndpoint REST API.

> **Note**
>
> To use the dual-stack endpoints for the data plane endpoints for your gateways, if you have already onboarded your gateways before December 1st, 2024, you may need to request IPv6 activation. For more information, see Using IPv6 with AWS IoT Wireless.

```
aws iotwireless get-service-endpoint
```

This command returns information about:

- The service type for which you want to get endpoint information about, which can be CUPS or LNS.
- The CUPS or LNS server trust certificate depending on the endpoint specified.
- Your data plane API endpoint in the following format:

```
account-specific-prefix.service.lorawan.aws-region.amazonaws.com
```

where *service* can be `cups` or `lns`.

The following table contains generic representations of the AWS Account-specific LNS endpoints for each Region that AWS IoT Core for LoRaWAN supports. In the **Endpoint** column, the *account-specific-prefix* from your Account-specific endpoint replaces *prefix* shown in the generic endpoint representation.

**LNS endpoints**

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | *prefix*.lns.lorawan.us-east-1.amazonaws.com | WSS |
| US West (Oregon) | us-west-2 | *prefix*.lns.lorawan.us-west-2.amazonaws.com | WSS |
| Europe (Ireland) | eu-west-1 | *prefix*.lns.lorawan.eu-west-1.amazonaws.com | WSS |
| Europe (Frankfurt) | eu-central-1 | *prefix*.lns.lorawan.eu-central-1.amazonaws.com | WSS |
| Asia Pacific (Tokyo) | ap-northeast-1 | | WSS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | *prefix*.lns.lorawan.ap-northeast-1.amazonaws.com | |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*.lns.lorawan.ap-southeast-2.amazonaws.com | WSS |
| South America (São Paulo) | sa-east-1 | *prefix*.lns.lorawan.sa-east-1.amazonaws.com | WSS |

The following table contains generic representations of the AWS Account-specific CUPS endpoints for each Region that AWS IoT Core supports. In the **Endpoint** column, the *account-specific-prefix* from your Account-specific endpoint replaces *prefix* shown in the generic endpoint representation.

**CUPS endpoints**

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | *prefix*.cups.lorawan.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | *prefix*.cups.lorawan.us-west-2.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | *prefix*.cups.lorawan.eu-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | *prefix*.cups.lorawan.eu-central-1.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*.cups.lorawan.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*.cups.lorawan.ap-southeast-2.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | *prefix*.cups.lorawan.sa-east-1.amazonaws.com | HTTPS |

For information about the actions supported by the AWS IoT Wireless - data plane endpoints, see AWS IoT Wireless API operations in the *AWS IoT Wireless API Reference*.

## Service quotas

The following tables describes the maximum number of transactions per second (TPS) that can be made to each action in the AWS IoT Wireless API, which includes AWS IoT Core for LoRaWAN and AWS IoT Core for Amazon Sidewalk.

**Contents**


- Wireless devices and device data quotas
- LoRaWAN gateways API throttling
- Profiles and destinations API throttling
- Sidewalk and logging API throttling
- Tagging and endpoint API throttling
- Device location and additional AWS IoT Wireless API limits


### Wireless devices and device data quotas

The following service quotas apply to LoRaWAN and Sidewalk devices. It also includes device metrics, and device data quotas, which are transmitted between wireless devices, gateways, and the cloud.

### AWS IoT Wireless devices API throttling

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| TPS limit for AssociateWirelessDeviceWithThing | TPS limit for AssociateWirelessDeviceWithThing | 10 | Yes |
| TPS limit for CreateNetworkAnalyzerConfiguration | TPS limit for CreateNetworkAnalyzerConfiguration | 10 | Yes |
| TPS limit for CreateWirelessDevice | TPS limit for CreateWirelessDevice | 10 | Yes |
| TPS limit for DeleteNetworkAnalyzerConfiguration | TPS limit for DeleteNetworkAnalyzerConfiguration | 10 | Yes |
| TPS limit for DeleteWirelessDevice | TPS limit for DeleteWirelessDevice | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for DisassociateWirelessDeviceFromThing | TPS limit for DisassociateWirelessDeviceFromThing | 10 | Yes |
| TPS limit for GetEventConfigurationByResourceTypes | TPS limit for GetEventConfigurationByResourceTypes | 10 | Yes |
| TPS limit for GetMetricConfiguration | TPS limit for GetMetricConfiguration | 10 | Yes |
| TPS limit for GetMetrics | TPS limit for GetMetrics | 10 | Yes |
| TPS limit for GetWirelessDevice | TPS limit for GetWirelessDevice | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for GetWirelessDeviceStatistics | TPS limit for GetWirelessDeviceStatistics | 10 | Yes |
| TPS limit for ListEventConfigurations | TPS limit for ListEventConfigurations | 10 | Yes |
| TPS limit for ListNetworkAnalyzerConfigurations | TPS limit for ListNetworkAnalyzerConfigurations | 10 | Yes |
| TPS limit for ListWirelessDevices | TPS limit for ListWirelessDevices | 10 | Yes |
| TPS limit for SendDataToWirelessDevice | TPS limit for SendDataToWirelessDevice | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for TestWirelessDevice | TPS limit for TestWirelessDevice | 10 | Yes |
| TPS limit for UpdateEventConfigurationByResourceTypes | TPS limit for UpdateEventConfigurationByResourceTypes | 10 | Yes |
| TPS limit for UpdateMetricConfiguration | TPS limit for UpdateMetricConfiguration | 10 | Yes |
| TPS limit for UpdateWirelessDevice | TPS limit for UpdateWirelessDevice | 10 | Yes |

## LoRaWAN gateways API throttling

This table describes the maximum TPS for APIs used with LoRaWAN gateways. The gateways route messages between LoRaWAN devices and AWS IoT Core for LoRaWAN.

### AWS IoT Wireless gateway API throttling

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| TPS limit for AssociateWirelessGatewayWithCertificate | TPS limit for AssociateWirelessGatewayWithCertificate | 10 | Yes |
| TPS limit for AssociateWirelessGatewayWithThing | TPS limit for AssociateWirelessGatewayWithThing | 10 | Yes |
| TPS limit for CreateWirelessGateway | TPS limit for CreateWirelessGateway | 10 | Yes |
| TPS limit for CreateWirelessGatewayTask | TPS limit for CreateWirelessGatewayTask | 10 | Yes |
| TPS limit for CreateWirelessGate | TPS limit for CreateWirelessGatewayTaskDefinition | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| wayTaskDefinition | | | |
| TPS limit for DeleteWirelessGateway | TPS limit for DeleteWirelessGateway | 10 | Yes |
| TPS limit for DeleteWirelessGatewayTask | TPS limit for DeleteWirelessGatewayTask | 10 | Yes |
| TPS limit for DeleteWirelessGatewayTaskDefinition | TPS limit for DeleteWirelessGatewayTaskDefinition | 10 | Yes |
| TPS limit for DisassociateWirelessGatewayFromCertificate | TPS limit for DisassociateWirelessGatewayFromCertificate | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for DisassociateWirelessGatewayFromThing | TPS limit for DisassociateWirelessGatewayFromThing | 10 | Yes |
| TPS limit for GetWirelessGateway | TPS limit for GetWirelessGateway | 10 | Yes |
| TPS limit for GetWirelessGatewayCertificate | TPS limit for GetWirelessGatewayCertificate | 10 | Yes |
| TPS limit for GetWirelessGatewayFirmwareInformation | TPS limit for GetWirelessGatewayFirmwareInformation | 10 | Yes |
| TPS limit for GetWirelessGatewayStatistics | TPS limit for GetWirelessGatewayStatistics | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for GetWirelessGatewayTask | TPS limit for GetWirelessGatewayTask | 10 | Yes |
| TPS limit for GetWirelessGateway TaskDefinition | TPS limit for GetWirelessGatewayTaskDefinition | 10 | Yes |
| TPS limit for ListWirelessGatewayTaskDefinitions | TPS limit for ListWirelessGatewayTaskDefinitions | 10 | Yes |
| TPS limit for ListWirelessGateways | TPS limit for ListWirelessGateways | 10 | Yes |
| TPS limit for UpdateWirelessGateway | TPS limit for UpdateWirelessGateway | 10 | Yes |

## Profiles and destinations API throttling

This table describes the maximum TPS for device profiles, service profiles, and destinations that can route messages to other AWS services.

### AWS IoT Wireless device profiles and destination API throttling

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| TPS limit for CreateDestination | TPS limit for CreateDestination | 10 | Yes |
| TPS limit for CreateDeviceProfile | TPS limit for CreateDeviceProfile | 10 | Yes |
| TPS limit for CreateServiceProfile | TPS limit for CreateServiceProfile | 10 | Yes |
| TPS limit for DeleteDestination | TPS limit for DeleteDestination | 10 | Yes |
| TPS limit for DeleteDeviceProfile | TPS limit for DeleteDeviceProfile | 10 | Yes |
| TPS limit for DeleteServiceProfile | TPS limit for DeleteServiceProfile | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for GetDestination | TPS limit for GetDestination | 10 | Yes |
| TPS limit for GetDeviceProfile | TPS limit for GetDeviceProfile | 10 | Yes |
| TPS limit for GetServiceProfile | TPS limit for GetServiceProfile | 10 | Yes |
| TPS limit for ListDestinations | TPS limit for ListDestinations | 10 | Yes |
| TPS limit for ListDeviceProfiles | TPS limit for ListDeviceProfiles | 10 | Yes |
| TPS limit for ListServiceProfiles | TPS limit for ListServiceProfiles | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for UpdateDestination | TPS limit for UpdateDestination | 10 | Yes |

## Sidewalk and logging API throttling

This table describes the maximum TPS for Amazon Sidewalk APIs and APIs that are used for log levels based on resource types.

### AWS IoT Wireless Sidewalk and logging API throttling

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for AssociateAwsAccountWithPartnerAccount | TPS limit for AssociateAwsAccountWithPartnerAccount | 10 | Yes |
| TPS limit for DeleteWirelessDeviceImportTask | TPS limit for DeleteWirelessDeviceImportTask | 10 | Yes |
| TPS limit for DeregisterWirelessDevice | TPS limit for DeregisterWirelessDevice | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for GetLogLevelsByResourceTypes | TPS limit for GetLogLevelsByResourceTypes | 10 | Yes |
| TPS limit for GetPartnerAccount | TPS limit for GetPartnerAccount | 10 | Yes |
| TPS limit for GetResourceLogLevel | TPS limit for GetResourceLogLevel | 10 | Yes |
| TPS limit for GetWirelessDeviceImportTask | TPS limit for GetWirelessDeviceImportTask | 10 | Yes |
| TPS limit for ListDevicesForWirelessDeviceImportTask | TPS limit for ListDevicesForWirelessDeviceImportTask | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for ListPartnerAccounts | TPS limit for ListPartnerAccounts | 10 | Yes |
| TPS limit for ListWirelessDeviceImportTasks | TPS limit for ListWirelessDeviceImportTasks | 10 | Yes |
| TPS limit for PutResourceLogLevel | TPS limit for PutResourceLogLevel | 10 | Yes |
| TPS limit for ResetAllResourceLogLevels | TPS limit for ResetAllResourceLogLevels | 10 | Yes |
| TPS limit for ResetResourceLogLevel | TPS limit for ResetResourceLogLevel | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for StartSingleWirelessDeviceImportTask | TPS limit for StartSingleWirelessDeviceImportTask | 10 | Yes |
| TPS limit for StartWirelessDeviceImportTask | TPS limit for StartWirelessDeviceImportTask | 10 | Yes |
| TPS limit for UpdateLogLevelsByResourceTypes | TPS limit for UpdateLogLevelsByResourceTypes | 10 | Yes |
| TPS limit for UpdatePartnerAccount | TPS limit for UpdatePartnerAccount | 10 | Yes |
| TPS limit for UpdateWirelessDeviceImportTask | TPS limit for UpdateWirelessDeviceImportTask | 10 | Yes |

## Tagging and endpoint API throttling

This table describes the maximum TPS for the `GetServiceEndpoint` API and APIs used for tagging resources.

### AWS IoT Wireless tagging and GetServiceEndpoint API throttling

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| TPS limit for GetServiceEndpoint | TPS limit for GetServiceEndpoint | 10 | Yes |
| TPS limit for ListTagsForResource | TPS limit for ListTagsForResource | 10 | Yes |
| TPS limit for TagResource | TPS limit for TagResource | 10 | Yes |
| TPS limit for UntagResource | TPS limit for UntagResource | 10 | Yes |

## Device location and additional AWS IoT Wireless API limits

This table describes the maximum TPS for other, additional AWS IoT Wireless API throttling including AWS IoT Core device location, the network analyzer feature, FUOTA, and multicast groups.

### AWS IoT Wireless limits and quotas

| Limit display name | Description | Default value | Adjustable |
|---|---|---|---|
| TPS limit for AssociateMulticastGroupWithFuotaTask | TPS limit for AssociateMulticastGroupWithFuotaTask | 10 | Yes |
| TPS limit for AssociateWirelessDeviceWithFuotaTask | TPS limit for AssociateWirelessDeviceWithFuotaTask | 10 | Yes |
| TPS limit for AssociateWirelessDeviceWithMulticastGroup | TPS limit for AssociateWirelessDeviceWithMulticastGroup | 10 | Yes |
| TPS limit for CancelMulticastGroupSession | TPS limit for CancelMulticastGroupSession | 10 | Yes |
| TPS limit for CreateFuotaTask | TPS limit for CreateFuotaTask | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for CreateMulticastGroup | TPS limit for CreateMulticastGroup | 10 | Yes |
| TPS limit for DeleteFuotaTask | TPS limit for DeleteFuotaTask | 10 | Yes |
| TPS limit for DeleteMulticastGroup | TPS limit for DeleteMulticastGroup | 10 | Yes |
| TPS limit for DeleteQueuedMessages | TPS limit for DeleteQueuedMessages | 10 | Yes |
| TPS limit for DisassociateAwsAccountFromPartnerAccount | TPS limit for DisassociateAwsAccountFromPartnerAccount | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for DisassociateMulticastGroupFromFuotaTask | TPS limit for DisassociateMulticastGroupFromFuotaTask | 10 | Yes |
| TPS limit for DisassociateWirelessDeviceFromFuotaTask | TPS limit for DisassociateWirelessDeviceFromFuotaTask | 10 | Yes |
| TPS limit for DisassociateWirelessDeviceFromMulticastGroup | TPS limit for DisassociateWirelessDeviceFromMulticastGroup | 10 | Yes |
| TPS limit for GetFuotaTask | TPS limit for GetFuotaTask | 10 | Yes |
| TPS limit for GetMulticastGroup | TPS limit for GetMulticastGroup | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for GetMulticastGroupSession | TPS limit for GetMulticastGroupSession | 10 | Yes |
| TPS limit for GetNetworkAnalyzerConfiguration | TPS limit for GetNetworkAnalyzerConfiguration | 10 | Yes |
| TPS limit for GetPosition | TPS limit for GetPosition | 10 | Yes |
| TPS limit for GetPositionConfiguration | TPS limit for GetPositionConfiguration | 10 | Yes |
| TPS limit for GetPositionEstimate | TPS limit for GetPositionEstimate | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for GetResourceEventConfiguration | TPS limit for GetResourceEventConfiguration | 10 | Yes |
| TPS limit for GetResourcePosition | TPS limit for GetResourcePosition | 10 | Yes |
| TPS limit for ListFuotaTasks | TPS limit for ListFuotaTasks | 10 | Yes |
| TPS limit for ListMulticastGroups | TPS limit for ListMulticastGroups | 10 | Yes |
| TPS limit for ListMulticastGroupsByFuotaTask | TPS limit for ListMulticastGroupsByFuotaTask | 10 | Yes |
| TPS limit for ListPositionConfigurations | TPS limit for ListPositionConfigurations | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for ListQueuedMessages | TPS limit for ListQueuedMessages | 10 | Yes |
| TPS limit for PutPositionConfiguration | TPS limit for PutPositionConfiguration | 10 | Yes |
| TPS limit for SendDataToMulticastGroup | TPS limit for SendDataToMulticastGroup | 10 | Yes |
| TPS limit for StartBulkAssociateWirelessDeviceWithMulticastGroup | TPS limit for StartBulkAssociateWirelessDeviceWithMulticastGroup | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for StartBulkDisassociateWirelessDeviceFromMulticastGroup | TPS limit for StartBulkDisassociateWirelessDeviceFromMulticastGroup | 10 | Yes |
| TPS limit for StartFuotaTask | TPS limit for StartFuotaTask | 10 | Yes |
| TPS limit for StartMulticastGroupSession | TPS limit for StartMulticastGroupSession | 10 | Yes |
| TPS limit for StartNetworkAnalyzerStream | TPS limit for StartNetworkAnalyzerStream | 10 | Yes |
| TPS limit for UpdateFuotaTask | TPS limit for UpdateFuotaTask | 10 | Yes |

| Limit display name | Description | Default value | Adjustable |
| --- | --- | --- | --- |
| TPS limit for UpdateMulticastGroup | TPS limit for UpdateMulticastGroup | 10 | Yes |
| TPS limit for UpdateNetworkAnalyzerConfiguration | TPS limit for UpdateNetworkAnalyzerConfiguration | 10 | Yes |
| TPS limit for UpdatePosition | TPS limit for UpdatePosition | 10 | Yes |
| TPS limit for UpdateResourceEventConfiguration | TPS limit for UpdateResourceEventConfiguration | 10 | Yes |
| TPS limit for UpdateResourcePosition | TPS limit for UpdateResourcePosition | 10 | Yes |

# AWS IoT Greengrass V1 endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Control Plane Operations

The following table contains AWS Region-specific endpoints that AWS IoT Greengrass supports for group management operations.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | greengrass.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | greengrass.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | greengrass.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | greengrass.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | greengrass.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | greengrass.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | greengrass.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | greengrass.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | greengrass.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | greengrass.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | greengrass.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | greengrass.us-gov-east-1.amazonaws.com<br>greengrass.us-gov-east-1.amazonaws.com<br>greengrass-ats.iot.us-gov-east-1.amazonaws.com<br>greengrass-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS<br>MQTT and HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | greengrass.us-gov-west-1.amazonaws.com<br>greengrass-ats.iot.us-gov-west-1.amazonaws.com<br>greengrass.us-gov-west-1.amazonaws.com | HTTPS<br>MQTT and HTTPS<br>HTTPS | |

## AWS IoT Device Operations

The following table contains AWS Region-specific Amazon Trust Services (ATS) endpoints for AWS IoT device management operations, such as shadow sync. This is a data plane API.

To look up your account-specific endpoint, use the aws iot describe-endpoint --endpoint-type iot:Data-ATS command.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | *prefix*-ats.iot.us-east-2.amazonaws.com | HTTPS, MQTT |
| US East (N. Virginia) | us-east-1 | *prefix*-ats.iot.us-east-1.amazonaws.com | HTTPS, MQTT |
| US West (Oregon) | us-west-2 | *prefix*-ats.iot.us-west-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Mumbai) | ap-south-1 | *prefix*-ats.iot.ap-south-1.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Seoul) | ap-northeast-2 | *prefix*-ats.iot.ap-northeast-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Singapore) | ap-southeast-1 | *prefix*-ats.iot.ap-southeast-1.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*-ats.iot.ap-southeast-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*-ats.iot.ap-northeast-1.amazonaws.com | HTTPS, MQTT |
| China (Beijing) | cn-north-1 | *prefix*.ats.iot.cn-north-1.amazonaws.com.cn | HTTPS, MQTT |
| Europe (Frankfurt) | eu-central-1 | *prefix*-ats.iot.eu-central-1.amazonaws.com | HTTPS, MQTT |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | *prefix*-ats.iot.eu-west-1.amazonaws.com | HTTPS, MQTT |
| Europe (London) | eu-west-2 | *prefix*-ats.iot.eu-west-2.amazonaws.com | HTTPS, MQTT |
| AWS GovCloud (US-West) | us-gov-west-1 | *prefix*-ats.iot.us-gov-west-1.amazonaws.com | HTTPS, MQTT |
| AWS GovCloud (US-East) | us-gov-east-1 | *prefix*-ats.iot.us-gov-east-1.amazonaws.com | HTTPS, MQTT |

**Note**

Legacy Verisign endpoints are currently supported for some Regions, but we recommend that you use ATS endpoints with ATS root certificate authority (CA) certificates. For more information, see Server Authentication in the *AWS IoT Developer Guide*.

## Discovery Operations

The following table contains AWS Region-specific ATS endpoints for device discovery operations using the AWS IoT Greengrass Discovery API. This is a data plane API.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | greengrass-ats.iot.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | greengrass-ats.iot.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | greengrass-ats.iot.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | greengrass-ats.iot.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | greengrass-ats.iot.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | greengrass-ats.iot.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | greengrass-ats.iot.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | greengrass-ats.iot.ap-northeast-1.amazonaws.com | HTTPS |
| China (Beijing) | cn-north-1 | greengrass.ats.iot.cn-north-1.amazonaws.com.cn | HTTPS |
| Europe (Frankfurt) | eu-central-1 | greengrass-ats.iot.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | greengrass-ats.iot.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | greengrass-ats.iot.eu-west-2.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | greengrass-ats.iot.us-gov-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | greengrass-ats.iot.us-gov-east-1.amazonaws.com | HTTPS |

**Note**

Legacy Verisign endpoints are currently supported for some Regions, but we recommend that you use ATS endpoints with ATS root CA certificates. For more information, see Server authentication in the *AWS IoT Developer Guide*.

## Supported Legacy Endpoints

We recommend that you use the ATS endpoints in the preceding tables with ATS root CA certificates. For backward compatibility, AWS IoT Greengrass currently supports legacy Verisign endpoints in the following AWS Regions. This support is expected to end in the future. For more information, see Server authentication in the *AWS IoT Developer Guide*.

When using legacy Verisign endpoints, you must use Verisign root CA certificates.

AWS IoT Device Operations (Legacy Endpoints)

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | *prefix*.iot.us-east-1.amazonaws.com | HTTPS, MQTT |
| US West (Oregon) | us-west-2 | *prefix*.iot.us-west-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*.iot.ap-southeast-2.amazonaws.com | HTTPS, MQTT |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*.iot.ap-northeast-1.amazonaws.com | HTTPS, MQTT |
| Europe (Frankfurt) | eu-central-1 | *prefix*.iot.eu-central-1.amazonaws.com | HTTPS, MQTT |
| Europe (Ireland) | eu-west-1 | *prefix*.iot.eu-west-1.amazonaws.com | HTTPS, MQTT |

To look up your account-specific legacy endpoint, use the aws iot describe-endpoint --endpoint-type iot:Data command.

Discovery Operations (Legacy Endpoints)

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | greengrass.iot.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | greengrass.iot.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | greengrass.iot.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | greengrass.iot.ap-northeast-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | greengrass.iot.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | greengrass.iot.eu-west-1.amazonaws.com | HTTPS |

# Service quotas

## AWS IoT Greengrass Cloud API

| Description | Default |
| --- | --- |
| Maximum number of AWS IoT devices per AWS IoT Greengrass group. | 2500 |
| Maximum number of Lambda functions per group. | 200 |
| Maximum number of resources per Lambda function. | 20 |
| Maximum number of resources per group. | 200 |
| Maximum number of transactions per second (TPS) on the AWS IoT Greengrass APIs. | See the section called “TPS”. |
| Maximum number of subscriptions per group. | 10000 |
| Maximum number of subscriptions that specify `Cloud` as the source per group. | 50 |
| Maximum length of a core thing name. | 124 bytes of UTF-8 encoded characters. |

### TPS

The default quota for the maximum number of transactions per second on the AWS IoT Greengrass APIs depends on the API and the AWS Region where AWS IoT Greengrass is used.

For most APIs and supported AWS Regions, the default quota is 30. Exceptions are noted in the following tables.

### API exceptions

| API | Default |
| --- | --- |
| CreateDeployment | 20 |

### AWS Region exceptions

| AWS Region | Default |
| --- | --- |
| China (Beijing) | 10 |
| AWS GovCloud (US-West) | 10 |
| AWS GovCloud (US-East) | 10 |

This quota applies per AWS account. For example, in the US East (N. Virginia) Region, each account has a default quota of 30 TPS. Each API (such as `CreateGroupVersion` or `ListFunctionDefinitions`) has a quota of 30 TPS. This includes control plane and data plane operations. Requests that exceed the account or API quotas are throttled. To request account and API quota increases, including quotas for specific APIs, contact your AWS Enterprise Support representative.

## AWS IoT Greengrass Core

| Description | Default |
| --- | --- |
| Maximum number of routing table entries that specify `Cloud` as the source. | 50 (matches AWS IoT subscription quota) |
| Maximum size of messages sent by an AWS IoT device. | 128 KB (matches AWS IoT message size quota) |
| Minimum message queue size in the Greengrass core router. | 256 KB |
| Maximum length of a topic string. | 256 bytes of UTF-8 encoded characters. |
| Maximum number of forward slashes (/) in a topic or topic filter. | 7 |
| Minimum disk space needed to run the Greengrass Core software. | 128 MB<br><br>400 MB when using OTA updates |

| Description | Default |
| --- | --- |
| Minimum RAM to run the Greengrass Core software. | 128 MB<br><br>198 MB when using stream manager |

The Greengrass Core software provides a service to detect the IP addresses of your Greengrass core devices. It sends this information to the AWS IoT Greengrass cloud service and allows AWS IoT devices to download the IP address of the Greengrass core they need to connect to.

Do not use this feature if any of the following is true:

- The IP address of a Greengrass core device changes frequently.
- The Greengrass core device is not always available to AWS IoT devices in its group.
- The Greengrass core has multiple IP addresses and an AWS IoT device is unable to reliably determine which address to use.
- Your organization's security policies don't allow you to send devices' IP addresses to the AWS Cloud.

# AWS IoT Greengrass V2 endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Control plane operations

The following table contains AWS Region-specific endpoints that AWS IoT Greengrass V2 supports for operations to manage components, devices, and deployments.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | greengrass.us-east-2.amazonaws.com<br>greengrass-fips.us-east-2.amazonaws.com<br>greengrass.us-east-2.api.aws<br>greengrass-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | greengrass.us-east-1.amazonaws.com<br>greengrass-fips.us-east-1.amazonaws.com<br>greengrass.us-east-1.api.aws<br>greengrass-fips.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | greengrass.us-west-2.amazonaws.com<br>greengrass-fips.us-west-2.amazonaws.com<br>greengrass.us-west-2.api.aws<br>greengrass-fips.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | greengrass.ap-southeast-5.amazonaws.com<br>greengrass.ap-southeast-5.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | greengrass.ap-south-1.amazonaws.com<br>greengrass.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | greengrass.ap-northeast-2.amazonaws.com<br>greengrass.ap-northeast-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | greengrass.ap-southeast-1.amazonaws.com<br>greengrass.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | greengrass.ap-southeast-2.amazonaws.com<br>greengrass.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | greengrass.ap-northeast-1.amazonaws.com<br>greengrass.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | greengrass.ca-central-1.amazonaws.com<br>greengrass-fips.ca-central-1.amazonaws.com<br>greengrass.ca-central-1.api.aws<br>greengrass-fips.ca-central-1.api.aws | HTTPS | |
| China (Beijing) | cn-north-1 | greengrass.cn-north-1.amazonaws.com.cn | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | greengrass.eu-central-1.amazonaws.com<br>greengrass.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | greengrass.eu-west-1.amazonaws.com<br>greengrass.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | greengrass.eu-west-2.amazonaws.com<br>greengrass.eu-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | greengrass.eu-south-2.amazonaws.com<br>greengrass.eu-south-2.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | greengrass.us-gov-west-1.amazonaws.com<br>greengrass.us-gov-west-1.api.aws | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | greengrass.us-gov-east-1.amazonaws.com<br>greengrass.us-gov-east-1.api.aws | HTTPS | |

## AWS IoT device operations

The following table contains AWS Region-specific Amazon Trust Services (ATS) endpoints for AWS IoT device management operations, such as shadow sync. This is a data plane API.

To look up your account-specific endpoint, use the aws iot describe-endpoint --endpoint-type iot:Data-ATS command.

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | *prefix*-ats.iot.us-east-2.amazonaws.com | HTTPS, MQTT |
| US East (N. Virginia) | us-east-1 | *prefix*-ats.iot.us-east-1.amazonaws.com | HTTPS, MQTT |
| US West (Oregon) | us-west-2 | *prefix*-ats.iot.us-west-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Malaysia) | ap-southeast-5 | *prefix*-ats.iot.ap-southeast-5.amazonaws.com | HTTPS, MQTT |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | *prefix*-ats.iot.ap-south-1.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Seoul) | ap-northeast-2 | *prefix*-ats.iot.ap-northeast-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Singapore) | ap-southeast-1 | *prefix*-ats.iot.ap-southeast-1.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*-ats.iot.ap-southeast-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*-ats.iot.ap-northeast-1.amazonaws.com | HTTPS, MQTT |
| Canada (Central) | ca-central-1 | *prefix*-ats.iot.ca-central-1.amazonaws.com | HTTPS, MQTT |
| China (Beijing) | cn-north-1 | *prefix*.ats.iot.cn-north-1.amazonaws.com.cn | HTTPS, MQTT |
| Europe (Frankfurt) | eu-central-1 | *prefix*-ats.iot.eu-central-1.amazonaws.com | HTTPS, MQTT |
| Europe (Ireland) | eu-west-1 | *prefix*-ats.iot.eu-west-1.amazonaws.com | HTTPS, MQTT |
| Europe (London) | eu-west-2 | *prefix*-ats.iot.eu-west-2.amazonaws.com | HTTPS, MQTT |
| Europe (Spain) | eu-south-2 | *prefix*-ats.iot.eu-south-2.amazonaws.com | HTTPS, MQTT |
| AWS GovCloud (US-West) | us-gov-west-1 | *prefix*-ats.iot.us-gov-west-1.amazonaws.com | HTTPS, MQTT |

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-East) | us-gov-east-1 | *prefix*-ats.iot.us-gov-east-1.amazonaws.com | HTTPS, MQTT |

> **Note**
>
> Legacy Verisign endpoints are currently supported for some Regions, but we recommend that you use ATS endpoints with ATS root certificate authority (CA) certificates. For more information, see Server Authentication in the *AWS IoT Developer Guide*.

## Data plane operations

AWS IoT Greengrass uses the AWS IoT Core Region-specific ATS endpoints for data plane operations, such as ResolveComponentCandidates. For a complete list, see AWS IoT Core - data plane endpoints. You must set `greengrassDataPlaneEndpoint` to **`iotdata`**. For more information, see AWS IoT Greengrass nucleus configuration.

We recommend you use the AWS IoT endpoints for data plane operations. For backwards compatibility, AWS IoT Greengrass supports the following legacy endpoints.

| Region name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | greengrass-ats.iot.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | greengrass-ats.iot.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | greengrass-ats.iot.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | greengrass-ats.iot.ap-southeast-5.amazonaws.com | HTTPS |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | greengrass-ats.iot.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | greengrass-ats.iot.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | greengrass-ats.iot.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | greengrass-ats.iot.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | greengrass-ats.iot.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | greengrass-ats.iot.ap-northeast-1.amazonaws.com | HTTPS |
| China (Beijing) | cn-north-1 | greengrass.ats.iot.cn-north-1.amazonaws.com.cn | HTTPS |
| Europe (Frankfurt) | eu-central-1 | greengrass-ats.iot.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | greengrass-ats.iot.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | greengrass-ats.iot.eu-west-2.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | greengrass-ats.iot.eu-south-2.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | greengrass-ats.iot.us-gov-west-1.amazonaws.com | HTTPS |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | greengrass-ats.iot.us-gov-east-1.amazonaws.com | HTTPS |

**Note**

Legacy Verisign endpoints are currently supported for some Regions, but we recommend that you use ATS endpoints with ATS root CA certificates. For more information, see Server authentication in the *AWS IoT Developer Guide*.

## Supported legacy endpoints

We recommend that you use the ATS endpoints in the preceding tables with ATS root CA certificates. For backward compatibility, AWS IoT Greengrass V2 currently supports legacy Verisign endpoints in the following AWS Regions. This support is expected to end in the future. For more information, see Server authentication in the *AWS IoT Developer Guide*.

When using legacy Verisign endpoints, you must use Verisign root CA certificates.

AWS IoT device operations (legacy endpoints)

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | *prefix*.iot.us-east-1.amazonaws.com | HTTPS, MQTT |
| US West (Oregon) | us-west-2 | *prefix*.iot.us-west-2.amazonaws.com | HTTPS, MQTT |
| Asia Pacific (Sydney) | ap-southeast-2 | *prefix*.iot.ap-southeast-2.amazonaws.com | HTTPS, MQTT |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | *prefix*.iot.ap-northeast-1.amazonaws.com | HTTPS, MQTT |
| Europe (Frankfurt) | eu-central-1 | *prefix*.iot.eu-central-1.amazonaws.com | HTTPS, MQTT |
| Europe (Ireland) | eu-west-1 | *prefix*.iot.eu-west-1.amazonaws.com | HTTPS, MQTT |

To look up your account-specific legacy endpoint, use the aws iot describe-endpoint --endpoint-type iot:Data command.

Data plane operations (legacy endpoints)

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | greengrass.iot.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | greengrass.iot.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | greengrass.iot.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | greengrass.iot.ap-northeast-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | greengrass.iot.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | greengrass.iot.eu-west-1.amazonaws.com | HTTPS |

# Service quotas

The following tables describe quotas in AWS IoT Greengrass V2. For more information about quotas and how to request quota increases, see AWS service quotas.

### Quotas for core devices

| Resource | Quota | Adjustable |
|---|---|---|
| Maximum length of a core device thing name | 124 bytes of UTF-8 encoded characters | No |

### Quotas for components

| Resource | Quota | Adjustable | Notes |
|---|---|---|---|
| Maximum number of components | 5,000 components per Region | No | |
| Maximum number of component versions | 5,000 versions per component per Region | No | |
| Maximum size of component recipe | 16 KB | No | |
| Maximum total size of component artifacts | 2 GB | No | This quota applies to the sum of all artifacts for a component. |

### Quotas for deployments

| Resource | Quota | Adjustable | Notes |
|---|---|---|---|
| Maximum size of deployment document for a thing | 7 KB | No | The deployment document includes the component |

| Resource | Quota | Adjustable | Notes |
| --- | --- | --- | --- |
| deployment (without large configuration support) | | | configurations, deployment configurations, and payload overhead. |
| Maximum size of deployment document for a thing group deployment (without large configuration support) | 31 KB | No | The deployment document includes the component configurations, deployment configurations, and payload overhead. |
| Maximum size of deployment document with large configuration support | 6 MB | No | The deployment document includes the component configurations, deployment configurations, and payload overhead. |

**Quotas for API operations**

| Resource | Quota | Adjustable | Notes |
| --- | --- | --- | --- |
| Request rate for CreateComponentVersion | 1 request per second per Region | No | |
| Request rate for other API operations | 30 requests per second per Region | No | This quota applies to the combination of API requests for all control plane operations. |

| Resource | Quota | Adjustable | Notes |
| --- | --- | --- | --- |
| | | | **Exceptions**<br>• China (Beijing) – 10 requests per second per Region<br>• AWS GovCloud (US-West) – 10 requests per second per Region<br>• AWS GovCloud (US-East) – 10 requests per second per Region |

# AWS IoT SiteWise endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The following sections describe the service endpoints for AWS IoT SiteWise.

**Topics**

- AWS IoT SiteWise – service endpoints
- AWS IoT SiteWise – control plane endpoints
- AWS IoT SiteWise – data plane endpoints

## AWS IoT SiteWise – service endpoints

A user can connect programmatically to AWS IoT SiteWise with an endpoint. Do not add a prefix to AWS SDK and AWS Command Line Interface commands because they add the necessary prefix.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | iotsitewise.us-east-2.amazonaws.com | HTTPS | |
| | | iotsitewise-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | iotsitewise.us-east-1.amazonaws.com | HTTPS | |
| | | iotsitewise-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | iotsitewise.us-west-2.amazonaws.com | HTTPS | |
| | | iotsitewise-fips.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | iotsitewise.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | iotsitewise.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | iotsitewise.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | iotsitewise.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | iotsitewise.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | iotsitewise.ca-central-1.amazonaws.com<br>iotsitewise-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | iotsitewise.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | iotsitewise.eu-west-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | iotsitewise.us-gov-west-1.amazonaws.com<br>iotsitewise-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## AWS IoT SiteWise – control plane endpoints

Use the api prefix for control plane APIs. For example, `api.iotsitewise.us-west-1.amazonaws.com`.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | api.iotsitewise.us-east-2.amazonaws.com<br>api.iotsitewise-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | api.iotsitewise.us-east-1.amazonaws.com<br>api.iotsitewise-fips.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | api.iotsitewise.us-west-2.amazonaws.com<br><br>api.iotsitewise-fips.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | api.iotsitewise.ap-south-1.amazonaws.com | | |
| Asia Pacific (Seoul) | ap-northeast-2 | api.iotsitewise.ap-northeast-2.amazonaws.com | | |
| Asia Pacific (Singapore) | ap-southeast-1 | api.iotsitewise.ap-southeast-1.amazonaws.com | | |
| Asia Pacific (Sydney) | ap-southeast-2 | api.iotsitewise.ap-southeast-2.amazonaws.com | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.iotsitewise.ap-northeast-1.amazonaws.com | | |
| Canada (Central) | ca-central-1 | api.iotsitewise.ca-central-1.amazonaws.com<br><br>api.iotsitewise-fips.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | api.iotsitewise.eu-central-1.amazonaws.com | | |
| Europe (Ireland) | eu-west-1 | api.iotsitewise.eu-west-1.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | api.iotsitewise.us-gov-west-1.amazonaws.com<br>api.iotsitewise-fips.us-gov-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | monitor.iotsitewise.us-east-2.amazonaws.com<br>monitor.iotsitewise-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | monitor.iotsitewise.us-east-1.amazonaws.com<br>monitor.iotsitewise-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | monitor.iotsitewise.us-west-2.amazonaws.com<br>monitor.iotsitewise-fips.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | monitor.iotsitewise.ap-south-1.amazonaws.com | | |
| Asia Pacific (Seoul) | ap-northeast-2 | monitor.iotsitewise.ap-northeast-2.amazonaws.com | | |
| Asia Pacific (Singapore) | ap-southeast-1 | monitor.iotsitewise.ap-southeast-1.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | monitor.iotsitewise.ap-southeast-2.amazonaws.com | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | monitor.iotsitewise.ap-northeast-1.amazonaws.com | | |
| Canada (Central) | ca-central-1 | monitor.iotsitewise.ca-central-1.amazonaws.com<br><br>monitor.iotsitewise-fips.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | monitor.iotsitewise.eu-central-1.amazonaws.com | | |
| Europe (Ireland) | eu-west-1 | monitor.iotsitewise.eu-west-1.amazonaws.com | | |
| AWS GovCloud (US-West) | us-gov-west-1 | monitor.iotsitewise.us-gov-west-1.amazonaws.com<br><br>monitor.iotsitewise-fips.us-gov-west-1.amazonaws.com | HTTPS | |

## AWS IoT SiteWise – data plane endpoints

Use the data prefix for data plane APIs. For example, `data.iotsitewise.us-west-1.amazonaws.com`.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | data.iotsitewise.us-east-2.amazonaws.com<br>data.iotsitewise-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | data.iotsitewise.us-east-1.amazonaws.com<br>data.iotsitewise-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | data.iotsitewise.us-west-2.amazonaws.com<br>data.iotsitewise-fips.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | data.iotsitewise.ap-south-1.amazonaws.com | | |
| Asia Pacific (Seoul) | ap-northeast-2 | data.iotsitewise.ap-northeast-2.amazonaws.com | | |
| Asia Pacific (Singapore) | ap-southeast-1 | data.iotsitewise.ap-southeast-1.amazonaws.com | | |
| Asia Pacific (Sydney) | ap-southeast-2 | data.iotsitewise.ap-southeast-2.amazonaws.com | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | data.iotsitewise.ap-northeast-1.amazonaws.com | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | data.iotsitewise.ca-central-1.amazonaws.com<br><br>data.iotsitewise-fips.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | data.iotsitewise.eu-central-1.amazonaws.com | | |
| Europe (Ireland) | eu-west-1 | data.iotsitewise.eu-west-1.amazonaws.com | | |
| AWS GovCloud (US-West) | us-gov-west-1 | data.iotsitewise.us-gov-west-1.amazonaws.com<br><br>data.iotsitewise-fips.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

The following tables describe quotas in AWS IoT SiteWise. For more information about quotas, and to request quota increases, see AWS service quotas.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Composite model depth | Each supported Region: 2 | Yes | The maximum depth of the composite model tree per asset model, including inline and component-based composite models. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Depth of asset model hierarchy tree | Each supported Region: 30 | Yes | The maximum asset model hierarchy tree depth. |
| Depth of interface hierarchy tree | Each supported Region: 30 | Yes | The maximum interface hierarchy tree depth |
| Enrichment-processed source video hours per Workspace | Each supported Region: 100,000 | Yes | The maximum cumulative duration of source video that can be enrichment-processed for a workspace, in hours. |
| Maximum computation nodes per pipeline | Each supported Region: 50 | Yes | The maximum number of computation nodes per pipeline. |
| Number of OPC UA sources per gateway | Each supported Region: 100 | No | The maximum number of OPC-UA sources per gateway. |
| Number of applications in each AWS Region for each AWS account | Each supported Region: 30 | Yes | The maximum number of applications that you can create in an AWS Region for an AWS account. |
| Number of applications in each workspace | Each supported Region: 1 | No | The maximum number of applications per workspace per Region per AWS account. |
| Number of asset models per Region per AWS account | Each supported Region: 10,000 | Yes | The maximum number of asset models per Region per AWS account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of asset models per hierarchy tree | Each supported Region: 100 | Yes | The maximum number of asset models per hierarchy tree. |
| Number of assets per asset model | Each supported Region: 10,000 | Yes | The maximum number of assets per asset model. |
| Number of child assets per parent asset | Each supported Region: 2,000 | Yes | The maximum number of child assets per parent asset. |
| Number of composite models per asset model | Each supported Region: 50 | Yes | The maximum number of composite models allowed on a single asset model. |
| Number of computation models per Region per AWS account | Each supported Region: 1,000 | Yes | The maximum number of computation models per Region per AWS account. |
| Number of concurrent in-flight DEEP searches | Each supported Region: 600 | Yes | The maximum number of concurrent in-flight DEEP searches per Region per AWS account. |
| Number of concurrent in-flight DEEP searches for Workspace | Each supported Region: 100 | Yes | The maximum number of concurrent in-flight DEEP searches for a workspace. |
| Number of concurrent pipeline executions per Workspace | Each supported Region: 200 | Yes | The maximum number of concurrent pipeline executions per workspace. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of concurrent running bulk import jobs | Each supported Region: 100 | Yes | The maximum number of bulk import jobs that can concurrently run |
| Number of dashboards per project | Each supported Region: 100 | Yes | The maximum number of dashboards per project. |
| Number of data binding values per binding value list | Each supported Region: 80 | No | The maximum number of data binding values per binding value list. |
| Number of data points per second per data quality per asset property | Each supported Region: 10 | No | The maximum number of data points with the same timestamp in seconds per data quality for each asset property. You can store up to this number of good-quality, uncertain-quality, and bad-quality data points for any given second for each asset property. |
| Number of data segments per dataset | Each supported Region: 10,000 | Yes | The maximum number of data segments per dataset per workspace per Region per AWS account. |
| Number of datasets per workspace | Each supported Region: 10,000 | Yes | The maximum number of datasets per workspace per Region per AWS account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of days between the timestamp in the past and today for Adaptive Ingestion in Bulk Import Job API | Each supported Region: 30 | No | The maximum number of days between a timestamp in the past and today date when using Adaptive Ingestion in Bulk Import Job API. |
| Number of event detection enrichment jobs for Workspace | Each supported Region: 120 | Yes | The maximum number of event detection enrichment jobs for a workspace. |
| Number of event detection enrichment jobs per account | Each supported Region: 240 | Yes | The maximum number of event detection enrichment jobs per Region per AWS account. |
| Number of functions per property formula expression | Each supported Region: 10 | No | The maximum number of functions that can be used in one property formula expression. |
| Number of gateways per Region per AWS account | Each supported Region: 100 | Yes | The maximum number of gateways per Region per AWS account. |
| Number of hierarchy definitions per asset model | Each supported Region: 30 | Yes | The maximum number of hierarchy definitions per asset model. |
| Number of hierarchy definitions per interface | Each supported Region: 20 | Yes | The maximum number of hierarchy definitions per interface |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of in-flight MP4 ingest jobs | Each supported Region: 100 | Yes | The maximum number of in-flight MP4 ingest jobs per Region per AWS account. |
| Number of in-flight MP4 ingest jobs for Workspace | Each supported Region: 50 | Yes | The maximum number of in-flight MP4 ingest jobs for a workspace. |
| Number of in-flight annotation ingest jobs | Each supported Region: 20 | Yes | The maximum number of in-flight annotation ingest jobs per Region per AWS account. |
| Number of in-flight annotation ingest jobs for Workspace | Each supported Region: 10 | Yes | The maximum number of in-flight annotation ingest jobs for a workspace. |
| Number of in-flight bulk import jobs per dataset | Each supported Region: 1 | No | The maximum number of in-flight bulk import jobs per dataset (all types combined) per Region per AWS account. |
| Number of in-flight export jobs | Each supported Region: 100 | Yes | The maximum number of in-flight export jobs per Region per AWS account. |
| Number of in-flight export jobs for Workspace | Each supported Region: 50 | Yes | The maximum number of in-flight export jobs for a workspace. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of in-flight queries | Each supported Region: 30 | Yes | The maximum number of in-flight queries per Region per AWS account. |
| Number of in-flight queries for Workspace | Each supported Region: 15 | Yes | The maximum number of in-flight queries for a workspace. |
| Number of in-flight telemetry ingest jobs | Each supported Region: 20 | Yes | The maximum number of in-flight telemetry ingest jobs per Region per AWS account. |
| Number of in-flight telemetry ingest jobs for Workspace | Each supported Region: 10 | Yes | The maximum number of in-flight telemetry ingest jobs for a workspace. |
| Number of interface per Region per AWS account | Each supported Region: 100 | Yes | The maximum number of interface per Region per AWS account |
| Number of interface per hierarchy tree | Each supported Region: 100 | Yes | The maximum number of interface per hierarchy tree |
| Number of metrics per dashboard visualization | Each supported Region: 5 | Yes | The maximum number of metrics per dashboard visualization. |
| Number of parent asset models per child asset model | Each supported Region: 10,000 | Yes | The maximum number of parent asset models per child asset model. |
| Number of pipelines per Workspace | Each supported Region: 200 | Yes | The maximum number of pipelines per workspace. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of portals per Region per AWS account | Each supported Region: 100 | Yes | The maximum number of portals per Region per AWS account. |
| Number of projects per portal | Each supported Region: 100 | Yes | The maximum number of projects per portal. |
| Number of properties at root of each asset model of type ASSET_MODEL | Each supported Region: 500 | Yes | This maximum number of assetModelProperties for each asset model. This count does not include composite ModelProperties. This limit also applies to any unique asset created from this asset model. |
| Number of properties per asset model | Each supported Region: 5,000 | Yes | The maximum number of properties of an asset model of type ASSET_MODEL or COMPONENT_MODEL. This number is determined by combining the properties of the root asset model and any included component-based or inline composite models. This limit also applies to any unique asset created from this asset model. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of properties per composite model | Each supported Region: 100 | Yes | The maximum number of properties allowed for composite models. Also, the maximum number of properties allowed for an asset model of type COMPONENT_MODEL. |
| Number of properties per interface | Each supported Region: 200 | Yes | The maximum number of properties in an interface |
| Number of properties that depend on a single property | Each supported Region: 30 | No | The maximum number of properties that directly or indirectly depend on a single property, as defined across all formula expressions. |
| Number of properties that depend on a single property within enforced asset model | Each supported Region: 10 | Yes | The maximum number of properties that directly or indirectly depend on a single property within enforced asset model |
| Number of properties that directly depend on a single property | Each supported Region: 20 | Yes | The maximum number of properties that directly depend on a single property, as defined across all formula expressions. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of property variables per property formula expression | Each supported Region: 10 | No | The maximum number of property variables that can be used in one property formula expression. |
| Number of results per GetInterpolatedAssetPropertyValues request | Each supported Region: 10 | Yes | The maximum number of results to return per paginated GetInterpolatedAssetPropertyValues request. |
| Number of root assets per project | Each supported Region: 1 | No | The maximum number of root assets associated per project. |
| Number of tasks per Workspace | Each supported Region: 800 | Yes | The maximum number of tasks per workspace. |
| Number of unique asset models that use the same component model | Each supported Region: 20 | Yes | The maximum number of unique asset models that have at least one component-based composite model that directly references a specific asset model of type COMPONENT_MODEL. |
| Number of unique asset models that use the same interface model | Each supported Region: 500 | Yes | The maximum number of unique asset models that use the same interface model |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of unique data streams in a single data bulk import job file | Each supported Region: 10,000 | No | The maximum number of unique data streams in a single data bulk import job file |
| Number of visualizations per dashboard | Each supported Region: 10 | Yes | The maximum number of visualizations per dashboard. |
| Number of workspaces per Region per AWS account | Each supported Region: 30 | Yes | The maximum number of workspaces per Region per AWS account. |
| Rate of BatchPutAssetPropertyValue entries ingested per asset property | Each supported Region: 10 | No | The maximum number of BatchPutAssetPropertyValue entries ingested per second per asset property per Region per AWS account. |
| Rate of GetAssetPropertyAggregates request and BatchGetAssetPropertyAggregates entry queries per asset property | Each supported Region: 50 | No | The maximum number of total GetAssetPropertyAggregates requests and BatchGetAssetPropertyAggregates entries for each asset property per second per Region per AWS Account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetAssetPropertyValue request and BatchGetAssetPropertyValue entry queries per asset property | Each supported Region: 500 | No | The maximum number of total GetAssetPropertyValue requests and BatchGetAssetPropertyValue entries for each asset property per second per Region per AWS Account. |
| Rate of GetAssetPropertyValueHistory request and BatchGetAssetPropertyValueHistory entry queries per asset property | Each supported Region: 30 | No | The maximum number of total GetAssetPropertyValueHistory requests and BatchGetAssetPropertyValueHistory entries for each asset property per second per Region per AWS Account. |
| Rate of GetInterpolatedAssetPropertyValues requests | Each supported Region: 500 | Yes | The maximum number of GetInterpolatedAssetPropertyValues requests per second that you can perform in this account in the current Region. |
| Rate of data points ingested | Each supported Region: 5,000 | Yes | The maximum number of timestamp-quality-value (TQV) data points ingested per second per Region per AWS account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of datapoints retrieved from GetAssetPropertyValueHistory and BatchGetAssetPropertyValueHistory | Each supported Region: 100 | Yes | The maximum byte rate (MB/second) of datapoints retrieved per second per Region per AWS account across GetAssetPropertyValueHistory and BatchGetAssetPropertyValueHistory. The response payload evaluated for this quota uses Timestamp-Quality-Value (TQV) fields for each datapoint and rounds the byte size for each API request to the next 4KB increment. |
| Request rate for AssociateAssets | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for AssociateAssets. |
| Request rate for AssociateTimeSeriesToAssetProperty | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for AssociateTimeSeriesToAssetProperty. |
| Request rate for BatchAssociateDataSegmentsToDataset for Workspace | Each supported Region: 10 | Yes | The maximum number of BatchAssociateDataSegmentsToDataset requests per second for a workspace |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for BatchDeleteDataset DataSegments for Workspace | Each supported Region: 10 | Yes | The maximum number of BatchDeleteDataset DataSegments requests per second for a workspace |
| Request rate for BatchDisassociateD ataSegmentsFromDataset for Workspace | Each supported Region: 10 | Yes | The maximum number of BatchDisassociateD ataSegmentsFromDat aset requests per second for a workspace |
| Request rate for BatchGetAssetPrope rtyAggregates | Each supported Region: 200 | Yes | The maximum number of requests per second per Region per AWS account for BatchGetAssetPrope rtyAggregates. |
| Request rate for BatchGetAssetPrope rtyValue | Each supported Region: 500 | Yes | The maximum number of requests per second per Region per AWS account for BatchGetAssetPrope rtyValue. |
| Request rate for BatchGetAssetPrope rtyValueHistory | Each supported Region: 200 | Yes | The maximum number of requests per second per Region per AWS account for BatchGetAssetPrope rtyValueHistory. |
| Request rate for BatchPutAssetPrope rtyValue | Each supported Region: 1,000 | Yes | The maximum number of requests per second per Region per AWS account for BatchPutAssetPrope rtyValue. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for CancelEnrichmentJob | Each supported Region: 10 | Yes | The maximum number of CancelEnrichmentJob requests per second per Region per AWS account. |
| Request rate for CancelQuery | Each supported Region: 100 | Yes | The maximum number of CancelQuery requests per second per Region per AWS account. |
| Request rate for CancelQuery for Workspace | Each supported Region: 50 | Yes | The maximum number of CancelQuery requests per second for a workspace. |
| Request rate for CreateAsset | Each supported Region: 50 | Yes | The maximum number of requests per second per Region per AWS account for CreateAsset. |
| Request rate for CreateAssetModel | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for CreateAssetModel. |
| Request rate for CreateAssetModelCompositeModel | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for CreateAssetModelCompositeModel. |
| Request rate for CreateBulkImportJob for Workspace | Each supported Region: 10 | Yes | The maximum number of CreateBulkImportJob requests per second for a workspace |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for CreateDataset for Workspace | Each supported Region: 10 | Yes | The maximum number of CreateDataset requests per second for a workspace |
| Request rate for CreateDatasetExportJob | Each supported Region: 10 | Yes | The maximum number of CreateDatasetExportJob requests per second per Region per AWS account. |
| Request rate for CreateDatasetExportJob for Workspace | Each supported Region: 5 | Yes | The maximum number of CreateDatasetExportJob requests per second for a workspace. |
| Request rate for CreateEnrichmentJob | Each supported Region: 10 | Yes | The maximum number of CreateEnrichmentJob requests per second per Region per AWS account. |
| Request rate for DeleteAsset | Each supported Region: 50 | Yes | The maximum number of requests per second per Region per AWS account for DeleteAsset. |
| Request rate for DeleteAssetModel | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for DeleteAssetModel. |
| Request rate for DeleteAssetModelCompositeModel | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for DeleteAssetModelCompositeModel. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for DeleteDataset for Workspace | Each supported Region: 10 | Yes | The maximum number of DeleteDataset requests per second for a workspace |
| Request rate for DeleteTimeSeries | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for DeleteTimeSeries. |
| Request rate for DeleteTimeSeries for Workspace | Each supported Region: 100 | Yes | The maximum number of DeleteTimeSeries requests per second for a workspace |
| Request rate for DescribeAction | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeAction. |
| Request rate for DescribeAsset | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeAsset. |
| Request rate for DescribeAssetCompositeModel | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeAssetCompositeModel. |
| Request rate for DescribeAssetModel | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeAssetModel. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for DescribeAssetModelCompositeModel | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeAssetModelCompositeModel. |
| Request rate for DescribeAssetProperty | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeAssetProperty. |
| Request rate for DescribeBulkImportJob for Workspace | Each supported Region: 50 | Yes | The maximum number of DescribeBulkImportJob requests per second for a workspace |
| Request rate for DescribeDataset for Workspace | Each supported Region: 30 | Yes | The maximum number of DescribeDataset requests per second for a workspace |
| Request rate for DescribeDatasetExportJob | Each supported Region: 50 | Yes | The maximum number of DescribeDatasetExportJob requests per second per Region per AWS account. |
| Request rate for DescribeDatasetExportJob for Workspace | Each supported Region: 25 | Yes | The maximum number of DescribeDatasetExportJob requests per second for a workspace. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for DescribeDefaultEncryptionConfiguration | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for DescribeDefaultEncryptionConfiguration. |
| Request rate for DescribeEnrichmentJob | Each supported Region: 100 | Yes | The maximum number of DescribeEnrichmentJob requests per second per Region per AWS account. |
| Request rate for DescribeLoggingOptions | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeLoggingOptions. |
| Request rate for DescribeLoggingOptions for Workspace | Each supported Region: 150 | Yes | The maximum number of DescribeLoggingOptions requests per second for a workspace |
| Request rate for DescribeQuery | Each supported Region: 100 | Yes | The maximum number of DescribeQuery requests per second per Region per AWS account. |
| Request rate for DescribeQuery for Workspace | Each supported Region: 50 | Yes | The maximum number of DescribeQuery requests per second for a workspace. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for DescribeStorageConfiguration | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for DescribeStorageConfiguration. |
| Request rate for DescribeTimeSeries | Each supported Region: 150 | Yes | The maximum number of requests per second per Region per AWS account for DescribeTimeSeries. |
| Request rate for DescribeTimeSeries for Workspace | Each supported Region: 150 | Yes | The maximum number of DescribeTimeSeries requests per second for a workspace |
| Request rate for DisassociateAssets | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for DisassociateAssets. |
| Request rate for DisassociateTimeSeriesFromAssetProperty | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for DisassociateTimeSeriesFromAssetProperty. |
| Request rate for ExecuteAction | Each supported Region: 20 | Yes | The maximum number of requests per second per Region per AWS account for ExecuteAction. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for GetAssetPropertyAggregates | Each supported Region: 1,000 | Yes | The maximum number of requests per second per Region per AWS account for GetAssetPropertyAggregates. |
| Request rate for GetAssetPropertyValue | Each supported Region: 1,000 | Yes | The maximum number of requests per second per Region per AWS account for GetAssetPropertyValue. |
| Request rate for GetAssetPropertyValueHistory | Each supported Region: 1,000 | Yes | The maximum number of requests per second per Region per AWS account for GetAssetPropertyValueHistory. |
| Request rate for GetCaptureData | Each supported Region: 200 | Yes | The maximum number of GetCaptureData requests per second per Region per AWS account. |
| Request rate for GetCaptureData for Workspace | Each supported Region: 100 | Yes | The maximum number of GetCaptureData requests per second for a workspace. |
| Request rate for GetQueryResults | Each supported Region: 100 | Yes | The maximum number of GetQueryResults requests per second per Region per AWS account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for GetQueryResults for Workspace | Each supported Region: 50 | Yes | The maximum number of GetQueryResults requests per second for a workspace. |
| Request rate for ListActions | Each supported Region: 30 | Yes | The maximum number of requests per second per Region per AWS account for ListActions. |
| Request rate for ListAssetModelCompositeModels | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for ListAssetModelCompositeModels. |
| Request rate for ListAssetModels | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for ListAssetModels. |
| Request rate for ListAssetRelationships | Each supported Region: 30 | Yes | The maximum number of requests per second per Region per AWS account for ListAssetRelationships. |
| Request rate for ListAssets | Each supported Region: 30 | Yes | The maximum number of requests per second per Region per AWS account for ListAssets. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for ListAssociatedAssets | Each supported Region: 30 | Yes | The maximum number of requests per second per Region per AWS account for ListAssociatedAssets. |
| Request rate for ListBulkImportJobs for Workspace | Each supported Region: 50 | Yes | The maximum number of ListBulkImportJobs requests per second for a workspace |
| Request rate for ListCompositionRelationships | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for ListCompositionRelationships. |
| Request rate for ListDatasetDataSegmentRelationships for Workspace | Each supported Region: 10 | Yes | The maximum number of ListDatasetDataSegmentRelationships requests per second for a workspace |
| Request rate for ListDatasetDataSegments for Workspace | Each supported Region: 30 | Yes | The maximum number of ListDatasetDataSegments requests per second for a workspace |
| Request rate for ListDatasetExportJobs | Each supported Region: 100 | Yes | The maximum number of ListDatasetExportJobs requests per second per Region per AWS account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for ListDatasetExportJobs for Workspace | Each supported Region: 50 | Yes | The maximum number of ListDatasetExportJobs requests per second for a workspace. |
| Request rate for ListDatasets for Workspace | Each supported Region: 10 | Yes | The maximum number of ListDatasets requests per second for a workspace |
| Request rate for ListEnrichmentJobs | Each supported Region: 50 | Yes | The maximum number of ListEnrichmentJobs requests per second per Region per AWS account. |
| Request rate for ListQueries | Each supported Region: 100 | Yes | The maximum number of ListQueries requests per second per Region per AWS account. |
| Request rate for ListQueries for Workspace | Each supported Region: 50 | Yes | The maximum number of ListQueries requests per second for a workspace. |
| Request rate for ListTagsForResource | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for ListTagsForResource. |
| Request rate for ListTimeSeries | Each supported Region: 30 | Yes | The maximum number of requests per second per Region per AWS account for ListTimeSeries. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for ListTimeSeries for Workspace | Each supported Region: 30 | Yes | The maximum number of ListTimeSeries requests per second for a workspace |
| Request rate for PutDefaultEncryptionConfiguration | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for PutDefaultEncryptionConfiguration. |
| Request rate for PutLoggingOptions | Each supported Region: 3 | Yes | The maximum number of PutLoggingOptions requests per second per Region per AWS account |
| Request rate for PutLoggingOptions for Workspace | Each supported Region: 3 | Yes | The maximum number of PutLoggingOptions requests per second for a workspace |
| Request rate for PutStorageConfiguration | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for PutStorageConfiguration. |
| Request rate for StartQuery | Each supported Region: 30 | Yes | The maximum number of StartQuery requests per second per Region per AWS account. |
| Request rate for StartQuery for Workspace | Each supported Region: 15 | Yes | The maximum number of StartQuery requests per second for a workspace. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for TagResource | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for TagResource. |
| Request rate for UntagResource | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for UntagResource. |
| Request rate for UpdateAsset | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for UpdateAsset. |
| Request rate for UpdateAssetModel | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for UpdateAssetModel. |
| Request rate for UpdateAssetModelCompositeModel | Each supported Region: 10 | Yes | The maximum number of requests per second per Region per AWS account for UpdateAssetModelCompositeModel. |
| Request rate for UpdateAssetProperty | Each supported Region: 100 | Yes | The maximum number of requests per second per Region per AWS account for UpdateAssetProperty. |
| Request rate for UpdateDataset for Workspace | Each supported Region: 10 | Yes | The maximum number of UpdateDataset requests per second for a workspace |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Size of the CSV file for Adaptive Ingestion | Each supported Region: 256 Megabits | No | The maximum size of the CSV file for Adaptive Ingestion. |
| Size of the uncompressed parquet file | Each supported Region: 256 Megabits | No | The maximum size of the uncompressed parquet file. |
| Size of the uncompressed parquet row group | Each supported Region: 64 Megabits | No | The maximum size of the uncompressed parquet row group. |

# AWS IoT TwinMaker endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

To connect programmatically to AWS IoT TwinMaker, use an endpoint. If you use an HTTP client, you need to prefix control plane and data plane APIs as follows. However, it is unnecessary to add a prefix to AWS SDK and AWS Command Line Interface commands because they automatically add the necessary prefix.

- Use the `api` prefix for control plane APIs. For example, `api.iottwinmaker.us-west-1.amazonaws.com`.
- Use the `data` prefix for data plane APIs. For example, `data.iottwinmaker.us-west-1.amazonaws.com`.

For more information, see Endpoints and quotas in the AWS IoT TwinMaker User Guide.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | iottwinmaker.us-east-1.amazonaws.com | HTTPS |
| | | data.iottwinmaker.us-east-1.amazonaws.com | HTTPS |
| | | iottwinmaker-fips.us-east-1.amazonaws.com | HTTPS |
| | | api.iottwinmaker.us-east-1.amazonaws.com | HTTPS |
| | | api.iottwinmaker-fips.us-east-1.amazonaws.com | HTTPS |
| | | data.iottwinmaker-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | iottwinmaker.us-west-2.amazonaws.com | HTTPS |
| | | data.iottwinmaker.us-west-2.amazonaws.com | HTTPS |
| | | iottwinmaker-fips.us-west-2.amazonaws.com | HTTPS |
| | | api.iottwinmaker.us-west-2.amazonaws.com | HTTPS |
| | | api.iottwinmaker-fips.us-west-2.amazonaws.com | HTTPS |
| | | data.iottwinmaker-fips.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | iottwinmaker.ap-south-1.amazonaws.com | HTTPS |
| | | data.iottwinmaker.ap-south-1.amazonaws.com | HTTPS |
| | | api.iottwinmaker.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | iottwinmaker.ap-northeast-2.amazonaws.com<br>data.iottwinmaker.ap-northeast-2.amazonaws.com<br>api.iottwinmaker.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | iottwinmaker.ap-southeast-1.amazonaws.com<br>data.iottwinmaker.ap-southeast-1.amazonaws.com<br>api.iottwinmaker.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | iottwinmaker.ap-southeast-2.amazonaws.com<br>data.iottwinmaker.ap-southeast-2.amazonaws.com<br>api.iottwinmaker.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | iottwinmaker.ap-northeast-1.amazonaws.com<br>data.iottwinmaker.ap-northeast-1.amazonaws.com<br>api.iottwinmaker.ap-northeast-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | iottwinmaker.eu-central-1.amazonaws.com<br>data.iottwinmaker.eu-central-1.amazonaws.com<br>api.iottwinmaker.eu-central-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | iottwinmaker.eu-west-1.amazonaws.com | HTTPS |
| | | data.iottwinmaker.eu-west-1.amazonaws.com | HTTPS |
| | | api.iottwinmaker.eu-west-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | iottwinmaker.us-gov-west-1.amazonaws.com | HTTPS |
| | | data.iottwinmaker.us-gov-west-1.amazonaws.com | HTTPS |
| | | iottwinmaker-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | api.iottwinmaker.us-gov-west-1.amazonaws.com | HTTPS |
| | | api.iottwinmaker-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | data.iottwinmaker-fips.us-gov-west-1.amazonaws.com | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Component type composition tree depth | Each supported Region: 5 | Yes | The maximum depth of a component type composition tree. |
| Component type composition tree size | Each supported Region: 50 | Yes | The maximum number of component types |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | in a component type composition tree. |
| Component types per workspace | Each supported Region: 1,000 | Yes | The maximum number of unique component types per workspace. |
| Components per entity | Each supported Region: 30 | Yes | The maximum number of components that can be defined on one entity. |
| Depth of component type hierarchy | Each supported Region: 10 | Yes | The maximum depth of the component type hierarchy tree. |
| Depth of entity hierarchy | Each supported Region: 30 | Yes | The maximum depth of the entity hierarchy tree. |
| Entities per workspace | Each supported Region: 50,000 | Yes | The maximum number of entities allowed per workspace. |
| Number of SiteWise export resources per job | Each supported Region: 5,000 | Yes | The maximum number of SiteWise export resources per job. |
| Number of SiteWise import resources per job | Each supported Region: 5,000 | Yes | The maximum number of SiteWise import resources per job. |
| Number of TwinMaker export resources per job | Each supported Region: 100,000 | Yes | The maximum number of TwinMaker export resources per job. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of TwinMaker import resources per job | Each supported Region: 100,000 | Yes | The maximum number of TwinMaker import resources per job. |
| Number of metadata transfer jobs in queue | Each supported Region: 10 | Yes | The maximum number of PENDING metadata transfer jobs. |
| Parent component types per child component type | Each supported Region: 10 | No | The maximum number of multi-inheritance parent component types or "extendsFrom" relationships one component type can have. |
| Properties per component type or component | Each supported Region: 800 | Yes | The maximum number of properties that can be defined on a component type or added to any given component instance. |
| Scenes per workspace | Each supported Region: 100 | Yes | The maximum number of scenes within a workspace. |
| Size of the metadata transfer job import file | Each supported Region: 100 | Yes | The maximum size of the imported file (in MB). |
| Tags per resource | Each supported Region: 50 | No | The maximum number of tags that can be placed on any resource with an ARN (like workspace, entity, component type, etc.). |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Workspaces in this account in the current Region | Each supported Region: 15 | Yes | The maximum number of workspaces in this account in the current Region. |

## AWS IoT TwinMaker API throttling limits

This topic describes the maximum number of transactions per second (TPS) that can be made to AWS IoT TwinMaker API actions per AWS account.

- **Request rate for Model and Scene API operations limit:** 10 TPS per AWS account.
- **Request rate for Data read and write API operations limit:** 100 TPS per AWS account.

> **Note**
>
> The TPS limits apply to all regions, and are not adjustable.

# Amazon Interactive Video Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

Amazon IVS uses an API for setting up and configuring IVS low-latency and real-time streaming applications. Amazon IVS Chat uses the main Chat API for setting up and managing chat rooms, and the Chat Messaging API for sending and receiving chat messages.

## IVS low-latency streaming endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ivs.us-east-1.amazonaws.com | HTTPS |
| | | ivs.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ivs.us-west-2.amazonaws.com | HTTPS |
| | | ivs.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ivs.ap-south-1.amazonaws.com | HTTPS |
| | | ivs.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ivs.ap-northeast-2.amazonaws.com | HTTPS |
| | | ivs.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ivs.ap-northeast-1.amazonaws.com | HTTPS |
| | | ivs.ap-northeast-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ivs.eu-central-1.amazonaws.com | HTTPS |
| | | ivs.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ivs.eu-west-1.amazonaws.com | HTTPS |
| | | ivs.eu-west-1.api.aws | HTTPS |

## IVS real-time streaming endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ivsrealtime.us-east-1.amazonaws.com | HTTPS |
| | | ivsrealtime.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ivsrealtime.us-west-2.amazonaws.com | HTTPS |
| | | ivsrealtime.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ivsrealtime.ap-south-1.amazonaws.com | HTTPS |
| | | ivsrealtime.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ivsrealtime.ap-northeast-2.amazonaws.com | HTTPS |
| | | ivsrealtime.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ivsrealtime.ap-northeast-1.amazonaws.com | HTTPS |
| | | ivsrealtime.ap-northeast-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ivsrealtime.eu-central-1.amazonaws.com | HTTPS |
| | | ivsrealtime.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ivsrealtime.eu-west-1.amazonaws.com | HTTPS |
| | | ivsrealtime.eu-west-1.api.aws | HTTPS |

## IVS Chat endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ivschat.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | ivschat.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ivschat.us-west-2.amazonaws.com | HTTPS |
| | | ivschat.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ivschat.ap-south-1.amazonaws.com | HTTPS |
| | | ivschat.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ivschat.ap-northeast-2.amazonaws.com | HTTPS |
| | | ivschat.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ivschat.ap-northeast-1.amazonaws.com | HTTPS |
| | | ivschat.ap-northeast-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ivschat.eu-central-1.amazonaws.com | HTTPS |
| | | ivschat.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ivschat.eu-west-1.amazonaws.com | HTTPS |
| | | ivschat.eu-west-1.api.aws | HTTPS |

## IVS Chat messaging endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | edge.ivschat.us-east-1.amazonaws.com | WSS |
| US West (Oregon) | us-west-2 | edge.ivschat.us-west-2.amazonaws.com | WSS |
| Asia Pacific (Mumbai) | ap-south-1 | edge.ivschat.ap-south-1.amazonaws.com | WSS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | edge.ivschat.ap-northeast-2.amazonaws.com | WSS |
| Asia Pacific (Tokyo) | ap-northeast-1 | edge.ivschat.ap-northeast-1.amazonaws.com | WSS |
| Europe (Frankfurt) | eu-central-1 | edge.ivschat.eu-central-1.amazonaws.com | WSS |
| Europe (Ireland) | eu-west-1 | edge.ivschat.eu-west-1.amazonaws.com | WSS |

## Service quotas

For more information, see Low-Latency Streaming Service Quotas in the *Amazon IVS Low-Latency Streaming User Guide* and Real-Time Streaming Service Quotas in the *Amazon IVS Real-Time Streaming User Guide*.

If there is a discrepancy, regard the *User Guides* as definitive.

### IVS low-latency streaming quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Ad configurations | 20 | No | Maximum number of ad configurations, per AWS Region. |
| Channels | 5,000 | Yes | Maximum number of channels, per AWS Region. |
| Concurrent streams | 100 | Yes | Maximum number of channels that can be streamed simultaneously, per AWS Region. If you exceed this threshold, the stream is rejected. |
| Concurrent views | 15,000 | Yes | Maximum number of views allowed to playback a live channel simultaneously, across |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | all channels, in an AWS Region. (A *view* is a unique viewing session which is actively downloading or playing video.) |
| Ingest bitrate (if channel `type` is `BASIC`) | 1.5 Mbps or 3.5 Mbps | No | Maximum bits per second that can be streamed to a channel whose `type` is `BASIC`.<br>• If input video quality is 480p or less, the default quota is 1.5 Mbps.<br>• If input video quality is more than 480p but less than 1080p, the default quota is 3.5 Mbps.<br><br>Warning: If you exceed this threshold, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |
| Ingest bitrate (if channel `type` is `STANDARD`, with single-track input) | 8.5 Mbps | No | Maximum bits per second that can be streamed to a channel whose `type` is `STANDARD` (the default), with single-track input. Warning: If you exceed this threshold, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Ingest bitrate (if channel `type` is `STANDARD`, with multitrack input up to `FULL_HD` resolution) | 15 Mbps | No | Maximum bits per second that can be streamed to a channel whose `type` is `STANDARD` (the default), whose `multitrackInputConfiguration.maximumResolution` is `SD`, `HD`, or `FULL_HD`, with multitrack input. Warning: If you exceed this threshold, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |
| Ingest bitrate (if channel `type` is `ADVANCED_HD`) | 8.5 Mbps | No | Maximum bits per second that can be streamed to a channel whose `type` is `ADVANCED_HD`. Warning: If you exceed this threshold, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |
| Ingest bitrate (if channel `type` is `ADVANCED_SD`) | 8.5 Mbps | No | Maximum bits per second that can be streamed to a channel whose `type` is `ADVANCED_SD`. Warning: If you exceed this threshold, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Ingest resolution (single-track input) | 1080p (2.1M total pixels, 1920 pixels/edge) | No | Maximum resolution in pixels that can be streamed to a channel (regardless of its `type`) with single-track input. There are two relevant thresholds: total pixels and pixels per edge. Warning: If you exceed either of these thresholds, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |
| Ingest resolution (if channel `type` is STANDARD, with multitrack input) | see Description | No | Maximum resolution in pixels that can be streamed to an individual track on a channel whose `type` is STANDARD, with multitrack input. The default varies depending on the value of `multitrackInputConfiguration.maximumResolution`:<br>• SD: 480p (0.4M total pixels, 864 pixels/edge)<br>• HD: 720p (0.9M total pixels, 1280 pixels/edge)<br>• FULL_HD: 1080p (2.1M total pixels, 1920 pixels/edge)><br><br>There are two relevant thresholds: total pixels and pixels per edge. Warning: If you exceed either of these thresholds, the stream probably will disconnect immediately. For details on channel `type`, see Channel Types in the *IVS Low-Latency Streaming API Reference*. |
| Metadata payload | 1 KB | No | Maximum size of a `PutMetadata` request payload (Amazon IVS API). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Playback authorization key pairs | 3 | No | Maximum number of playback authorization key pairs, per AWS Region. |
| Playback restriction policies | 3 | No | Maximum number of playback restriction policies, per AWS Region. |
| Playback restriction policy countries | 200 | No | Maximum size of the `allowedCountries` list in a playback restriction policy; that is, the maximum number of countries per policy. |
| Playback restriction policy origins | 5 | No | Maximum size of the `allowedOrigins` list in a playback restriction policy; that is, the maximum number of origins per policy. |
| Playback restriction policy origin length | 256 | No | Maximum size (in characters) of an entry in the `allowedOrigins` list in a playback restriction policy. |
| Playback token size | 2 KB | No | Maximum size of the entire JSON web token (JWT) used to initiate playback. |
| PutMetadata rate per channel | 5 | No | Maximum `PutMetadata` transactions per second per channel. |
| Recording configurations | 20 | Yes | Maximum number of recording configurations, per AWS Region. |
| Stream key | 1 | No | Maximum number of stream keys, per channel. |
| Stream takeovers | 100 | Yes | Maximum number of stream takeovers, per stream. |

## IVS real-time streaming quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Composition destinations | 2 | No | Maximum number of Destination objects in a Composition resource. |
| Composition: max duration | 24 | No | Maximum amount of time a composition can exist, in hours. |
| Compositions | 20 | Yes | Maximum concurrent Composition resources per account. |
| Compositions per stage | 5 | Yes | Maximum concurrent Composition resources per stage. |
| Concurrent participant replications | 5 | No | Maximum number of concurrent replications per participant across all stages in an AWS Region. |
| Concurrent publishers | 1,000 | Yes | Maximum number of participants who can be publishing across all stages in an AWS Region. |
| Concurrent subscriptions | 20,000 | Yes | Maximum number of simultaneous publisher-to-subscriber connections across all stages in an AWS Region. |
| EncoderConfigurations | 20 | Yes | Maximum number of EncoderConfiguration resources per account. |
| IngestConfigurations | 100 | Yes | Maximum number of IngestConfiguration resources per account. |
| Participant download bitrate | 8.5 Mbps | No | Maximum aggregate download bitrate across all of a participant's subscriptions. |
| Participant publish bitrate | 8.5 Mbps | No | Maximum bits per second that can be streamed to a stage. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Participant publish or subscribe duration | 24 | No | Maximum length of time a participant can publish or remain subscribed to a stage, in hours. |
| Participant publish resolution | 720p | No | Maximum resolution of video published by participants. |
| PublicKeys | 3 | No | Maximum number of public keys, per AWS Region. |
| Stage participants (publishers) | 12 | No | Maximum number of participants who can be publishing to a stage at once. |
| Stage participants (subscribers) | 10,000 | Yes | Maximum number of participants who can be subscribing to a stage at once. |
| Stages | 1,000 | Yes | Maximum number of stages, per AWS Region. |
| StorageConfigurations | 5 | Yes | Maximum number of StorageConfiguration resources per account. |

## IVS Chat quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent chat connections | 50,000 | Yes | Maximum number of concurrent chat connections per account, across all your rooms in an AWS Region. |
| Logging configurations | 10 | Yes | The maximum number of logging configurations that can be created per account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Message review handler timeout period | 200 | No | Timeout period in milliseconds for all your message review handlers in the current AWS Region. If this is exceeded, the message is allowed or denied depending on the value of the `fallbackResult` field you configured for the message review handler. |
| Rate of DeleteMessage requests across all your rooms | 100 | Yes | Maximum number of `DeleteMessage` requests that can be made per second across all your rooms. The requests can come from either the Amazon IVS Chat API or the Amazon IVS Chat Messaging API (WebSocket). |
| Rate of DisconnectUser requests across all your rooms | 100 | Yes | Maximum number of `DisconnectUser` requests that can be made per second across all your rooms. The requests can come from either the Amazon IVS Chat API or the Amazon IVS Chat Messaging API (WebSocket). |
| Rate of messaging requests per connection | 10 | No | Maximum number of messaging requests per second that a chat connection can make. |
| Rate of SendMessage requests across all your rooms | 1,000 | Yes | Maximum number of `SendMessage` requests that can be made per second across all your rooms. These requests come from the Amazon IVS Chat Messaging API (WebSocket). |
| Rate of SendMessage requests per room | 100 | No | Maximum number of `SendMessage` requests that can be made per second for any one of your rooms. This is configurable with the `maximumMessageRatePerSecond` field of Create Room and Update Room. These requests come from the Amazon IVS Chat Messaging API (WebSocket). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rooms | 50,000 | Yes | Maximum number of chat rooms per account, per AWS Region. |

# Amazon Kendra endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Amazon Kendra

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | kendra.us-east-2.amazonaws.com | HTTPS | |
| | | kendra-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | kendra.us-east-1.amazonaws.com | HTTPS | |
| | | kendra-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | kendra.us-west-2.amazonaws.com | HTTPS | |
| | | kendra-fips.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | kendra.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | kendra.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | kendra.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | kendra.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | kendra.ca-central-1.amazonaws.com<br>kendra-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | kendra.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | kendra.eu-west-2.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | kendra.us-gov-west-1.amazonaws.com<br>kendra-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

**Note**

Amazon Kendra GenAI Enterprise Edition indices are only available in the US East (N. Virginia), US West (Oregon), Europe (Ireland), Asia Pacific (Sydney), regions.

## Amazon Kendra Intelligent Ranking

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | kendra-ranking.us-east-2.api.aws | HTTPS |
| | | kendra-ranking-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | kendra-ranking.us-east-1.api.aws | HTTPS |
| | | kendra-ranking-fips.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | kendra-ranking.us-west-2.api.aws | HTTPS |
| | | kendra-ranking-fips.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | kendra-ranking.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | kendra-ranking.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | kendra-ranking.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | kendra-ranking.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | kendra-ranking.ca-central-1.api.aws<br><br>kendra-ranking-fips.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | kendra-ranking.eu-west-1.api.aws | HTTPS | |

# Service quotas

## Amazon Kendra

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Characters in query text | Each supported Region: 1,000 | Yes | The maximum number of characters in query text. |
| Collapsed results to expand | Each supported Region: 20 | Yes | The maximum number of collapsed result items to expand |
| Data sources (developer edition) | Each supported Region: 5 | No | The maximum number of data sources per developer edition index. |
| Data sources (enterprise edition) | Each supported Region: 50 | Yes | The maximum number of data sources per enterprise edition index. |
| Developer edition indexes | Each supported Region: 10 | Yes | The maximum number of developer edition indexes that can be created per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Enterprise edition indexes | Each supported Region: 10 | Yes | The maximum number of enterprise edition indexes that can be created per account. |
| Expanded results per collapsed result item | Each supported Region: 5 | Yes | The maximum number of expanded results per item |
| Extracted text size | Each supported Region: 5 Megabytes | No | The maximum extracted text size, in MB, from a file that can be ingested. |
| FAQs | Each supported Region: 30 | Yes | The maximum number of FAQs that can be created per index. |
| Featured documents per featured results set | Each supported Region: 4 | Yes | The maximum number of featured documents per featured results set. |
| Featured results sets per index | Each supported Region: 50 | Yes | The maximum number of featured results sets per index. |
| File size | Each supported Region: 50 Megabytes | Yes | The maximum file size, in MB, that can be ingested. |
| Ingestion attributes string list size | Each supported Region: 10 | Yes | The maximum string list size per ingestion attribute. |
| Items in a query suggestions block list | Each supported Region: 20,000 | Yes | The maximum number of items in a query suggestions block list. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Query attributes user group list size | Each supported Region: 10 | Yes | The maximum user group list size per query attribute. |
| Query capacity units | Each supported Region: 100 | Yes | The maximum number of query capacity units per index. |
| Query suggestions block list file size | Each supported Region: 2 Megabytes | Yes | The maximum query suggestions block list file size in MB. |
| Query suggestions block lists | Each supported Region: 1 | No | The maximum number of query suggestions block lists per index. |
| Query suggestions max additional response attributes in API | Each supported Region: 5 | Yes | The maximum number of additional fields/attributes to include in the response for query suggestions per GetQuerySuggestions API call. |
| Query suggestions max suggestable attributes in API | Each supported Region: 10 | Yes | The maximum number of fields/attributes for query suggestions per GetQuerySuggestions API call |
| Query suggestions returned in API | Each supported Region: 10 | Yes | The maximum number of query suggestions returned in a GetQuerySuggestions API call. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Sort configurations | Each supported Region: 3 | Yes | The maximum number of sort configurations |
| Spell correction query suggestions | Each supported Region: 1 | Yes | The maximum number of spell-corrected query suggestions to return in a Query API call. |
| Storage capacity units | Each supported Region: 100 | Yes | The maximum number of storage capacity units per index. |
| Synonym rules per thesaurus | Each supported Region: 10,000 | Yes | The maximum number of synonym rules per thesaurus. |
| Synonyms per term | Each supported Region: 10 | No | The maximum number of synonyms per term in all thesauri in a index. |
| Thesauri | Each supported Region: 1 | No | The maximum number of thesauri per index. |
| Thesaurus file size | Each supported Region: 5 Megabytes | Yes | The maximum thesaurus file size in MB. |

### Amazon Kendra Intelligent Ranking

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Maximum documents per group per Rescore request | Each supported Region: 3 | No | The maximum number of documents per group per Rescore request. |
| Maximum rescore capacity units per execution plan | Each supported Region: 1,000 | No | The maximum number of rescore capacity units per execution plan. |
| Maximum tokens per document body | Each supported Region: 200 | No | The maximum number of tokens in the body of a single document. |
| Maximum tokens per document title | Each supported Region: 100 | No | The maximum number of tokens in the title of a single document. |
| Maximum unique documents per Rescore request | Each supported Region: 25 | No | The maximum number of unique documents (or unique groups) per Rescore request. |
| Rescore capacity units per account | Each supported Region: 50 | Yes | The maximum number of rescore execution plans per account. |

## Amazon Keyspaces (for Apache Cassandra) endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | cassandra.us-east-2.amazonaws.com<br>cassandra.us-east-2.api.aws | HTTPS and TLS<br>HTTPS and TLS |
| US East (N. Virginia) | us-east-1 | cassandra.us-east-1.amazonaws.com<br>cassandra-fips.us-east-1.api.aws<br>cassandra-fips.us-east-1.amazonaws.com<br>cassandra.us-east-1.api.aws | HTTPS and TLS<br>HTTPS and TLS<br>HTTPS and TLS<br>HTTPS and TLS |
| US West (N. California) | us-west-1 | cassandra.us-west-1.amazonaws.com<br>cassandra.us-west-1.api.aws | HTTPS and TLS<br>HTTPS and TLS |
| US West (Oregon) | us-west-2 | cassandra.us-west-2.amazonaws.com<br>cassandra-fips.us-west-2.api.aws<br>cassandra-fips.us-west-2.amazonaws.com<br>cassandra.us-west-2.api.aws | HTTPS and TLS<br>HTTPS and TLS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | | HTTPS and TLS<br><br>HTTPS and TLS | |
| Africa (Cape Town) | af-south-1 | cassandra.af-south-1.amazonaws.com<br><br>cassandra.af-south-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Hong Kong) | ap-east-1 | cassandra.ap-east-1.amazonaws.com<br><br>cassandra.ap-east-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | cassandra.ap-southeast-5.amazonaws.com<br><br>cassandra.ap-southeast-5.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Mumbai) | ap-south-1 | cassandra.ap-south-1.amazonaws.com<br><br>cassandra.ap-south-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Seoul) | ap-northeast-2 | cassandra.ap-northeast-2.amazonaws.com<br><br>cassandra.ap-northeast-2.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | cassandra.ap-southeast-1.amazonaws.com<br><br>cassandra.ap-southeast-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Sydney) | ap-southeast-2 | cassandra.ap-southeast-2.amazonaws.com<br><br>cassandra.ap-southeast-2.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Thailand) | ap-southeast-7 | cassandra.ap-southeast-7.amazonaws.com<br><br>cassandra.ap-southeast-7.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | cassandra.ap-northeast-1.amazonaws.com<br><br>cassandra.ap-northeast-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Canada (Central) | ca-central-1 | cassandra.ca-central-1.amazonaws.com<br><br>cassandra.ca-central-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| Canada West (Calgary) | ca-west-1 | cassandra.ca-west-1.amazonaws.com<br><br>cassandra.ca-west-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | cassandra.eu-central-1.amazonaws.com<br>cassandra.eu-central-1.api.aws | HTTPS and TLS<br>HTTPS and TLS | |
| Europe (Ireland) | eu-west-1 | cassandra.eu-west-1.amazonaws.com<br>cassandra.eu-west-1.api.aws | HTTPS and TLS<br>HTTPS and TLS | |
| Europe (London) | eu-west-2 | cassandra.eu-west-2.amazonaws.com<br>cassandra.eu-west-2.api.aws | HTTPS and TLS<br>HTTPS and TLS | |
| Europe (Paris) | eu-west-3 | cassandra.eu-west-3.amazonaws.com<br>cassandra.eu-west-3.api.aws | HTTPS and TLS<br>HTTPS and TLS | |
| Europe (Stockholm) | eu-north-1 | cassandra.eu-north-1.amazonaws.com<br>cassandra.eu-north-1.api.aws | HTTPS and TLS<br>HTTPS and TLS | |
| Middle East (Bahrain) | me-south-1 | cassandra.me-south-1.amazonaws.com<br>cassandra.me-south-1.api.aws | HTTPS and TLS<br>HTTPS and TLS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | cassandra.me-central-1.amazonaws.com<br><br>cassandra.me-central-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| South America (São Paulo) | sa-east-1 | cassandra.sa-east-1.amazonaws.com<br><br>cassandra.sa-east-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| AWS GovCloud (US-East) | us-gov-east-1 | cassandra.us-gov-east-1.amazonaws.com<br><br>cassandra.us-gov-east-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |
| AWS GovCloud (US-West) | us-gov-west-1 | cassandra.us-gov-west-1.amazonaws.com<br><br>cassandra.us-gov-west-1.api.aws | HTTPS and TLS<br><br>HTTPS and TLS | |

For the following AWS Regions, FIPS endpoints are available.

| Region Name | Region | FIPS Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | cassandra-fips.us-east-1.amazonaws.com | TLS |
| US West (Oregon) | us-west-2 | cassandra-fips.us-west-2.amazonaws.com | TLS |

# Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Account-level read throughput quota (Provisioned mode) | Each supported Region: 80,000 | Yes | The maximum number of aggregate read capacity units (RCUs) allocated for the account per region; applicable only for tables in provisioned read/write capacity mode. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Account-level write throughput quota (Provisioned mode) | Each supported Region: 80,000 | Yes | The maximum number of aggregate write capacity units (WCUs) allocated for the account per region; applicable only for tables in provisioned read/write capacity mode. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Concurrent DDL operations | Each supported Region: 50 | No | The maximum number of concurrent DDL operations allowed per region. For more information, see https://docs.aws.amazon.com |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | /keyspaces/latest/devguide/quotas.html |
| Keyspaces per region | Each supported Region: 256 | Yes | The maximum number of keyspaces that can be created per region. |
| Max Schema size | Each supported Region: 358,400 Bytes | No | The maximum size for a table schema. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Max UDT name length | Each supported Region: 48 | Yes | The maximum number of characters in the UDT name. |
| Max amount of data restored using Point-in-time Recovery (PITR) | Each supported Region: 5 Terabytes | Yes | The maximum size of data that can be restored using PITR within 24 hours. |
| Max amount of data restored using add Region operations | Each supported Region: 10 Terabytes | Yes | The maximum size of data that can be concurrently restored by add Region operations. |
| Max amount of direct child UDTs per UDT | Each supported Region: 10 | Yes | The maximum number of child UDTs supported for a UDT. |
| Max amount of direct parent UDTs per UDT | Each supported Region: 10 | Yes | The maximum number of parent UDTs supported for a UDT. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max clustering key size | Each supported Region: 850 Bytes | No | The maximum combined size of all clustering columns. Up to 4 bytes of additional storage are added to the raw size of each clustering column for metadata. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Max concurrent table restores using Point-in-time Recovery (PITR) | Each supported Region: 4 | Yes | The maximum number of concurrent table restores using PITR per subscriber is 4 |
| Max number of UDTs per AWS Region | Each supported Region: 256 | Yes | The maximum number of UDTs that can be created per region across all keyspaces. |
| Max number of UDTs per table | Each supported Region: 50 | Yes | The maximum number of UDTs that a table can reference. |
| Max number of tables per UDT | Each supported Region: 100 | Yes | The Maximum number of tables that can refer to a particular UDT. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max partition key size | Each supported Region: 2,048 Bytes | No | The maximum size of the compound partition key. Up to 3 bytes of additional storage are added to the raw size of each column included in the partition key for metadata. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Max row size | Each supported Region: 1 Megabytes | No | The maximum size of a row, excluding static column data. For details see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Max static data per logical partition | Each supported Region: 1 Megabytes | No | The maximum aggregate size of static data in a logical partition. For details see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Max subqueries per IN SELECT statement | Each supported Region: 100 | Yes | The maximum number of subqueries you can use for the IN keyword in a SELECT statement. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Max write capacity for a table using change-data-capture (CDC) | Each supported Region: 40,000 | Yes | The maximum write capacity for Amazon Keyspaces tables with CDC enabled. This default quota only applies to tables in provisioned read/write capacity mode. |
| Table-level read throughput quota | Each supported Region: 40,000 | Yes | The maximum read throughput (RCUs & RRUs per second) that can be allocated to a table in the region. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |
| Table-level write throughput quota | Each supported Region: 40,000 | Yes | The maximum write throughput (WCUs & WRUs per second) that can be allocated to a table per region. For more information, see https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Tables per region | Each supported Region: 256 | Yes | The maximum number of tables that can be created per region across all keyspaces. |

For more information, see Quotas for Amazon Keyspaces (for Apache Cassandra) in the *Amazon Keyspaces (for Apache Cassandra) Developer Guide*.

# AWS Key Management Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | kms.us-east-2.amazonaws.com | HTTPS | |
| | | kms-fips.us-east-2.api.aws | HTTPS | |
| | | kms.us-east-2.api.aws | HTTPS | |
| | | kms-fips.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | kms.us-east-1.amazonaws.com | HTTPS |
| | | kms-fips.us-east-1.api.aws | HTTPS |
| | | kms-fips.us-east-1.amazonaws.com | HTTPS |
| | | kms.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | kms.us-west-1.amazonaws.com | HTTPS |
| | | kms.us-west-1.api.aws | HTTPS |
| | | kms-fips.us-west-1.amazonaws.com | HTTPS |
| | | kms-fips.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | kms.us-west-2.amazonaws.com | HTTPS |
| | | kms-fips.us-west-2.amazonaws.com | HTTPS |
| | | kms.us-west-2.api.aws | HTTPS |
| | | kms-fips.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | kms.af-south-1.amazonaws.com | HTTPS |
| | | kms.af-south-1.api.aws | HTTPS |
| | | kms-fips.af-south-1.api.aws | HTTPS |
| | | kms-fips.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | kms.ap-east-1.amazonaws.com | HTTPS |
| | | kms.ap-east-1.api.aws | HTTPS |
| | | kms-fips.ap-east-1.api.aws | HTTPS |
| | | kms-fips.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | kms.ap-south-2.amazonaws.com | HTTPS |
| | | kms-fips.ap-south-2.amazonaws.com | HTTPS |
| | | kms.ap-south-2.api.aws | HTTPS |
| | | kms-fips.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | kms.ap-southeast-3.amazonaws.com | HTTPS |
| | | kms.ap-southeast-3.api.aws | HTTPS |
| | | kms-fips.ap-southeast-3.api.aws | HTTPS |
| | | kms-fips.ap-southeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | kms.ap-southeast-5.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-5.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-5.api.aws | HTTPS |
| | | kms.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | kms.ap-southeast-4.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-4.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-4.api.aws | HTTPS |
| | | kms.ap-southeast-4.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | kms.ap-south-1.amazonaws.com | HTTPS |
| | | kms.ap-south-1.api.aws | HTTPS |
| | | kms-fips.ap-south-1.amazonaws.com | HTTPS |
| | | kms-fips.ap-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | kms.ap-southeast-6.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-6.api.aws | HTTPS |
| | | kms-fips.ap-southeast-6.amazonaws.com | HTTPS |
| | | kms.ap-southeast-6.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | kms.ap-northeast-3.amazonaws.com | HTTPS |
| | | kms-fips.ap-northeast-3.amazonaws.com | HTTPS |
| | | kms-fips.ap-northeast-3.api.aws | HTTPS |
| | | kms.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | kms.ap-northeast-2.amazonaws.com | HTTPS |
| | | kms.ap-northeast-2.api.aws | HTTPS |
| | | kms-fips.ap-northeast-2.api.aws | HTTPS |
| | | kms-fips.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | kms.ap-southeast-1.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-1.amazonaws.com | HTTPS |
| | | kms.ap-southeast-1.api.aws | HTTPS |
| | | kms-fips.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | kms.ap-southeast-2.amazonaws.com | HTTPS |
| | | kms-fips.ap-southeast-2.api.aws | HTTPS |
| | | kms-fips.ap-southeast-2.amazonaws.com | HTTPS |
| | | kms.ap-southeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | kms.ap-east-2.amazonaws.com | HTTPS |
| | | kms-fips.ap-east-2.api.aws | HTTPS |
| | | kms-fips.ap-east-2.amazonaws.com | HTTPS |
| | | kms.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | kms.ap-southeast-7.amazonaws.com | HTTPS |
| | | kms.ap-southeast-7.api.aws | HTTPS |
| | | kms-fips.ap-southeast-7.api.aws | HTTPS |
| | | kms-fips.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | kms.ap-northeast-1.amazonaws.com | HTTPS |
| | | kms-fips.ap-northeast-1.api.aws | HTTPS |
| | | kms.ap-northeast-1.api.aws | HTTPS |
| | | kms-fips.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | kms.ca-central-1.amazonaws.com | HTTPS |
| | | kms.ca-central-1.api.aws | HTTPS |
| | | kms-fips.ca-central-1.api.aws | HTTPS |
| | | kms-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | kms.ca-west-1.amazonaws.com | HTTPS |
| | | kms-fips.ca-west-1.api.aws | HTTPS |
| | | kms-fips.ca-west-1.amazonaws.com | HTTPS |
| | | kms.ca-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | kms.eu-central-1.amazonaws.com | HTTPS |
| | | kms.eu-central-1.api.aws | HTTPS |
| | | kms-fips.eu-central-1.api.aws | HTTPS |
| | | kms-fips.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | kms.eu-west-1.amazonaws.com | HTTPS |
| | | kms-fips.eu-west-1.api.aws | HTTPS |
| | | kms-fips.eu-west-1.amazonaws.com | HTTPS |
| | | kms.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | kms.eu-west-2.amazonaws.com | HTTPS |
| | | kms-fips.eu-west-2.amazonaws.com | HTTPS |
| | | kms.eu-west-2.api.aws | HTTPS |
| | | kms-fips.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | kms.eu-south-1.amazonaws.com | HTTPS |
| | | kms-fips.eu-south-1.amazonaws.com | HTTPS |
| | | kms-fips.eu-south-1.api.aws | HTTPS |
| | | kms.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | kms.eu-west-3.amazonaws.com | HTTPS |
| | | kms-fips.eu-west-3.amazonaws.com | HTTPS |
| | | kms-fips.eu-west-3.api.aws | HTTPS |
| | | kms.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | kms.eu-south-2.amazonaws.com | HTTPS | |
| | | kms-fips.eu-south-2.amazonaws.com | HTTPS | |
| | | kms.eu-south-2.api.aws | HTTPS | |
| | | kms-fips.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | kms.eu-north-1.amazonaws.com | HTTPS | |
| | | kms-fips.eu-north-1.amazonaws.com | HTTPS | |
| | | kms.eu-north-1.api.aws | HTTPS | |
| | | kms-fips.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | kms.eu-central-2.amazonaws.com | HTTPS | |
| | | kms-fips.eu-central-2.amazonaws.com | HTTPS | |
| | | kms-fips.eu-central-2.api.aws | HTTPS | |
| | | kms.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | kms.il-central-1.amazonaws.com | HTTPS | |
| | | kms-fips.il-central-1.amazonaws.com | HTTPS | |
| | | kms-fips.il-central-1.api.aws | HTTPS | |
| | | kms.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | kms.mx-central-1.amazonaws.com | HTTPS | |
| | | kms.mx-central-1.api.aws | HTTPS | |
| | | kms-fips.mx-central-1.amazonaws.com | HTTPS | |
| | | kms-fips.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | kms.me-south-1.amazonaws.com | HTTPS |
| | | kms-fips.me-south-1.api.aws | HTTPS |
| | | kms.me-south-1.api.aws | HTTPS |
| | | kms-fips.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | kms.me-central-1.amazonaws.com | HTTPS |
| | | kms-fips.me-central-1.amazonaws.com | HTTPS |
| | | kms-fips.me-central-1.api.aws | HTTPS |
| | | kms.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | kms.sa-east-1.amazonaws.com | HTTPS |
| | | kms-fips.sa-east-1.api.aws | HTTPS |
| | | kms.sa-east-1.api.aws | HTTPS |
| | | kms-fips.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | kms.us-gov-east-1.amazonaws.com | HTTPS |
| | | kms-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | kms-fips.us-gov-east-1.api.aws | HTTPS |
| | | kms.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | kms.us-gov-west-1.amazonaws.com | HTTPS |
| | | kms.us-gov-west-1.api.aws | HTTPS |
| | | kms-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | kms-fips.us-gov-west-1.api.aws | HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Aliases per CMK | Each supported Region: 50 | Yes | The maximum number of customer-created aliases per CMK permitted in each AWS Region of this AWS account. Aliases that AWS creates in your account with the aws/ prefix do not count against this quota. An alias is a friendly name for a customer master key (CMK). Each alias is associated with one CMK, but a CMK can have multiple aliases. |
| CancelKeyDeletion request rate | Each supported Region: 5 per second | Yes | Maximum CancelKeyDeletion requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ConnectCustomKeyStore request rate | Each supported Region: 5 per second | Yes | Maximum ConnectCustomKeyStore requests per second. When you reach this quota, KMS rejects requests for this operation for |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | the remainder of the interval. |
| CreateAlias request rate | Each supported Region: 5 per second | Yes | Maximum CreateAlias requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| CreateCustomKeyStore request rate | Each supported Region: 5 per second | Yes | Maximum CreateCustomKeyStore requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| CreateGrant request rate | Each supported Region: 50 per second | Yes | Maximum CreateGrant requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateKey request rate | Each supported Region: 5 per second | Yes | Maximum CreateKey requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| Cryptographic operations (ECC & SM2) request rate | Each supported Region: 1,000 per second | Yes | Maximum DeriveSharedSecret, Sign, and Verify requests with ECC KMS keys and Decrypt, DeriveSharedSecret, Encrypt, Sign, and Verify requests with SM2 (China Regions only) KMS keys per second. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cryptographic operations (ML-DSA) request rate | Each supported Region: 1,000 per second | Yes | Maximum requests for cryptographic operations with ML-DSA CMKs per second. This shared quota applies to Sign and Verify requests using ML-DSA CMKs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |
| Cryptographic operations (RSA) request rate | Each supported Region: 1,000 per second | Yes | Maximum requests for cryptographic operations with RSA CMKs per second. This shared quota applies to Encrypt, Decrypt, ReEncrypt, Sign, and Verify requests using RSA CMKs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Cryptographic operations (symmetric) request rate | us-east-1: 100,000 per second<br><br>us-east-2: 20,000 per second<br><br>us-west-2: 100,000 per second<br><br>ap-northeast-1: 20,000 per second<br><br>ap-southeast-1: 20,000 per second<br><br>ap-southeast-2: 20,000 per second<br><br>eu-central-1: 20,000 per second<br><br>eu-south-2: 20,000 per second<br><br>eu-west-1: 100,000 per second | Yes | Maximum requests for cryptographic operations with a symmetric CMK per second. This shared quota applies to Decrypt, Encrypt, GenerateDataKey, GenerateDataKeyWithoutPlaintext, GenerateMac, GenerateRandom, ReEncrypt, and VerifyMac requests. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | eu-west-2: 20,000 per second<br><br>Each of the other supported Regions: 10,000 per second | | |
| Custom Key Stores | Each supported Region: 10 | Yes | The maximum number of custom key stores permitted in each AWS Region of this AWS account. This quota applies to the total number of custom key stores, including AWS CloudHSM key stores and external key stores, regardless of their connection state. |
| Customer Master Keys (CMKs) | Each supported Region: 100,000 | Yes | The maximum number of customer managed CMKs permitted in each AWS Region of this AWS account. This quota does not apply to AWS managed CMKs. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteAlias request rate | Each supported Region: 15 per second | Yes | Maximum DeleteAlias requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| DeleteCustomKeyStore request rate | Each supported Region: 5 per second | Yes | Maximum DeleteCustomKeyStore requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| DeleteImportedKeyMaterial request rate | Each supported Region: 15 per second | Yes | Maximum DeleteImportedKeyMaterial requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DescribeCustomKeyStores request rate | Each supported Region: 5 per second | Yes | Maximum DescribeCustomKeyStores requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| DescribeKey request rate | Each supported Region: 2,000 per second | Yes | Maximum DescribeKey requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| DisableKey request rate | Each supported Region: 5 per second | Yes | Maximum DisableKey requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DisableKeyRotation request rate | Each supported Region: 5 per second | Yes | Maximum DisableKeyRotation requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| DisconnectCustomKeyStore request rate | Each supported Region: 5 per second | Yes | Maximum DisconnectCustomKeyStore requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| EnableKey request rate | Each supported Region: 5 per second | Yes | Maximum EnableKey requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| EnableKeyRotation request rate | Each supported Region: 15 per second | Yes | Maximum EnableKeyRotation requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| GenerateDataKeyPair (ECC_NIST_EDWARDS25519) request rate | Each supported Region: 100 per second | Yes | Maximum requests per second to generate ECC_NIST_EDWARDS25519 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for ECC_NIST_EDWARDS25519 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GenerateDataKeyPair (ECC_NIST_P256) request rate | Each supported Region: 100 per second | Yes | Maximum requests per second to generate ECC_NIST_P256 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for ECC_NIST_P256 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |
| GenerateDataKeyPair (ECC_NIST_P384) request rate | Each supported Region: 100 per second | Yes | Maximum requests per second to generate ECC_NIST_P384 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for ECC_NIST_P384 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GenerateDataKeyPair (ECC_NIST_P521) request rate | Each supported Region: 100 per second | Yes | Maximum requests per second to generate ECC_NIST_P521 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for ECC_NIST_P521 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |
| GenerateDataKeyPair (ECC_SECG_P256K1) request rate | Each supported Region: 100 per second | Yes | Maximum requests per second to generate ECC_SECG_P256K1 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for ECC_SECG_P256K1 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GenerateDataKeyPair (RSA_2048) request rate | Each supported Region: 20 per second | Yes | Maximum requests per second to generate RSA_2048 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for RSA_2048 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |
| GenerateDataKeyPair (RSA_3072) request rate | Each supported Region: 4 per second | Yes | Maximum requests per second to generate RSA_3072 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for RSA_3072 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GenerateDataKeyPair (RSA_4096) request rate | Each supported Region: 1 per second | Yes | Maximum requests per second to generate RSA_4096 data key pairs. This shared quota applies to GenerateDataKeyPair and GenerateDataKeyPairWithoutPlaintext requests for RSA_4096 data key pairs. When you reach this quota, KMS rejects this type of request for the remainder of the interval. |
| GetKeyPolicy request rate | Each supported Region: 1,000 per second | Yes | Maximum GetKeyPolicy requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| GetKeyRotationStatus request rate | Each supported Region: 1,000 per second | Yes | Maximum GetKeyRotationStatus requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetParametersForImport request rate | Each supported Region: 1 per second | Yes | Maximum GetParametersForImport requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| GetPublicKey request rate | Each supported Region: 2,000 per second | Yes | Maximum GetPublicKey requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| Grants per CMK | Each supported Region: 50,000 | Yes | The maximum number of grants permitted for each customer managed CMK. This quota includes grants created by AWS services, but it does not apply to AWS managed CMKs. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ImportKeyMaterial request rate | Each supported Region: 15 per second | Yes | Maximum ImportKeyMaterial requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ListAliases request rate | Each supported Region: 500 per second | Yes | Maximum ListAliases requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ListGrants request rate | Each supported Region: 100 per second | Yes | Maximum ListGrants requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListKeyPolicies request rate | Each supported Region: 100 per second | Yes | Maximum ListKeyPolicies requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ListKeyRotations request rate | Each supported Region: 100 per second | Yes | Maximum ListKeyRotations requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ListKeys request rate | Each supported Region: 500 per second | Yes | Maximum ListKeys requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListResourceTags request rate | Each supported Region: 2,000 per second | Yes | Maximum ListResourceTags requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ListRetirableGrants request rate | Each supported Region: 100 per second | Yes | Maximum ListRetirableGrants requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| On-demand rotations per CMK | Each supported Region: 25 | No | The maximum number of on-demand rotations permitted for each customer managed key. This quota does not apply to AWS managed keys. |
| PutKeyPolicy request rate | Each supported Region: 15 per second | Yes | Maximum PutKeyPolicy requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ReplicateKey request rate | Each supported Region: 5 per second | Yes | Maximum ReplicateKey requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| RetireGrant request rate | Each supported Region: 50 per second | Yes | Maximum RetireGrant requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| RevokeGrant request rate | Each supported Region: 50 per second | Yes | Maximum RevokeGrant requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| RotateKeyOnDemand request rate | Each supported Region: 5 per second | No | Maximum RotateKeyOnDemand requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| ScheduleKeyDeletion request rate | Each supported Region: 15 per second | Yes | Maximum ScheduleKeyDeletion requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| TagResource request rate | Each supported Region: 10 per second | Yes | Maximum TagResource requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UntagResource request rate | Each supported Region: 5 per second | Yes | Maximum UntagResource requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| UpdateAlias request rate | Each supported Region: 5 per second | Yes | Maximum UpdateAlias requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| UpdateCustomKeyStore request rate | Each supported Region: 5 per second | Yes | Maximum UpdateCustomKeyStore requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateKeyDescription request rate | Each supported Region: 5 per second | Yes | Maximum UpdateKeyDescription requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |
| UpdatePrimaryRegion request rate | Each supported Region: 5 per second | Yes | Maximum UpdatePrimaryRegion requests per second. When you reach this quota, KMS rejects requests for this operation for the remainder of the interval. |

## Amazon Kinesis Data Streams endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | kinesis.us-east-2.amazonaws.com | HTTPS | |
| | | kinesis-fips.us-east-2.api.aws | HTTPS | |
| | | kinesis-fips.us-east-2.amazonaws.com | HTTPS | |
| | | kinesis.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | kinesis.us-east-1.amazonaws.com | HTTPS | |
| | | kinesis-fips.us-east-1.api.aws | HTTPS | |
| | | kinesis-fips.us-east-1.amazonaws.com | HTTPS | |
| | | kinesis.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | kinesis.us-west-1.amazonaws.com | HTTPS | |
| | | kinesis-fips.us-west-1.api.aws | HTTPS | |
| | | kinesis-fips.us-west-1.amazonaws.com | HTTPS | |
| | | kinesis.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | kinesis.us-west-2.amazonaws.com | HTTPS | |
| | | kinesis-fips.us-west-2.api.aws | HTTPS | |
| | | kinesis-fips.us-west-2.amazonaws.com | HTTPS | |
| | | kinesis.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | kinesis.af-south-1.amazonaws.com | HTTPS | |
| | | kinesis.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | kinesis.ap-east-1.amazonaws.com<br>kinesis.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | kinesis.ap-south-2.amazonaws.com<br>kinesis.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | kinesis.ap-southeast-3.amazonaws.com<br>kinesis.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | kinesis.ap-southeast-5.amazonaws.com<br>kinesis.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | kinesis.ap-southeast-4.amazonaws.com<br>kinesis.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | kinesis.ap-south-1.amazonaws.com<br>kinesis.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | kinesis.ap-southeast-6.amazonaws.com<br>kinesis.ap-southeast-6.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | kinesis.ap-northeast-3.amazonaws.com<br>kinesis.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | kinesis.ap-northeast-2.amazonaws.com<br>kinesis.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | kinesis.ap-southeast-1.amazonaws.com<br>kinesis.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | kinesis.ap-southeast-2.amazonaws.com<br>kinesis.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | kinesis.ap-east-2.amazonaws.com<br>kinesis.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | kinesis.ap-southeast-7.amazonaws.com<br>kinesis.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | kinesis.ap-northeast-1.amazonaws.com<br>kinesis.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | kinesis.ca-central-1.amazonaws.com<br>kinesis.ca-central-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | kinesis.ca-west-1.amazonaws.com | HTTPS | |
| | | kinesis.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | kinesis.eu-central-1.amazonaws.com | HTTPS | |
| | | kinesis.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | kinesis.eu-west-1.amazonaws.com | HTTPS | |
| | | kinesis.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | kinesis.eu-west-2.amazonaws.com | HTTPS | |
| | | kinesis.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | kinesis.eu-south-1.amazonaws.com | HTTPS | |
| | | kinesis.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | kinesis.eu-west-3.amazonaws.com | HTTPS | |
| | | kinesis.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | kinesis.eu-south-2.amazonaws.com | HTTPS | |
| | | kinesis.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | kinesis.eu-north-1.amazonaws.com | HTTPS | |
| | | kinesis.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | kinesis.eu-central-2.amazonaws.com | HTTPS | |
| | | kinesis.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | kinesis.il-central-1.amazonaws.com<br>kinesis.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | kinesis.mx-central-1.amazonaws.com<br>kinesis.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | kinesis.me-south-1.amazonaws.com<br>kinesis.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | kinesis.me-central-1.amazonaws.com<br>kinesis.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | kinesis.sa-east-1.amazonaws.com<br>kinesis.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | kinesis.us-gov-east-1.amazonaws.com<br>kinesis.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | kinesis.us-gov-west-1.amazonaws.com<br>kinesis.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Shards per Region | us-east-1: 20,000<br>us-east-2: 6,000<br>us-west-2: 20,000<br>ap-northeast-1: 6,000<br>ap-south-1: 6,000<br>ap-southeast-1: 6,000<br>ap-southeast-2: 6,000<br>eu-central-1: 6,000<br>eu-west-1: 20,000<br>eu-west-3: 6,000<br>Each of the other supported Regions: 1,000 | Yes | The maximum number of shards that you can create in this account in the current Region. |

If you're using strict security policies that require explicit allowlisting of service endpoints, you must also allowlist the following account-level endpoints to ensure Kinesis Data Streams actions succeed.

- **Control-plane APIs**: `*.control-kinesis.<Region>.amazonaws.com and *.control-kinesis.<Region>.api.aws`

- **Data-plane APIs**: `*.data-kinesis.<Region>.amazonaws.com and *.data-kinesis.<Region>.api.aws`

The `.api.aws` endpoints are dual stack endpoints that accept IPv6 requests.

For more information, see Amazon Kinesis Data Streams Quotas in the *Amazon Kinesis Data Streams Developer Guide*.

# Amazon Kinesis Video Streams endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | kinesisvideo.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | kinesisvideo.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | kinesisvideo.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | kinesisvideo.af-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | kinesisvideo.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | kinesisvideo.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | kinesisvideo.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | kinesisvideo.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | kinesisvideo.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | kinesisvideo.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | kinesisvideo.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | kinesisvideo.ca-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | kinesisvideo.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | kinesisvideo.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | kinesisvideo.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | kinesisvideo.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | kinesisvideo.eu-south-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | kinesisvideo.il-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | kinesisvideo.me-south-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | kinesisvideo.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | kinesisvideo-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | kinesisvideo-fips.us-gov-west-1.amazonaws.com | HTTPS | |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ConnectAsMaster GO_AWAY message grace period | Each supported Region: 60 Seconds | No | The maximum duration of the grace period that follows the received GO_AWAY message, after which the master connection is terminated in this account in the current Region. |
| ConnectAsMaster connection duration | Each supported Region: 3,600 Seconds | No | The maximum duration of a master connection per signaling channel in this account in the current Region. |
| ConnectAsMaster connections per signaling channel | Each supported Region: 1 | No | The maximum number of master connections per signaling channel in this account in the current Region. |
| ConnectAsMaster idle connection timeout | Each supported Region: 600 Seconds | No | The maximum duration of an idle master connection before the connection times out in this account in the current Region. |
| ConnectAsViewer GO_AWAY message grace period | Each supported Region: 60 Seconds | No | The maximum duration of the grace period that follows the received GO_AWAY message, after which the viewer |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | connection is terminated in this account in the current Region. |
| ConnectAsViewer connection duration | Each supported Region: 3,600 Seconds | No | The maximum duration of a viewer connection per signaling channel in this account in the current Region. |
| ConnectAsViewer connections per signaling channel | Each supported Region: 10 | Yes | The maximum number of viewer connections per signaling channel in this account in the current Region. |
| ConnectAsViewer idle connection timeout | Each supported Region: 600 Seconds | No | The maximum duration of an idle viewer connection before the connection times out in this account in the current Region. |
| GetClip file size | Each supported Region: 100 Megabytes | No | The maximum size of a video file in megabytes retrieved by the GetClip API in this AWS account and Region. |
| GetClip fragments | Each supported Region: 200 | No | The maximum number of fragments that can be contained in a video file retrieved by the GetClip API in this AWS account and Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetDASHManifestPlaylist fragments | Each supported Region: 5,000 | No | The maximum number of fragments per GetDASHManifestPlaylist request per playlist in this account in the current Region. |
| GetHLSMediaPlaylist fragments | Each supported Region: 5,000 | No | The maximum number of fragments per GetHLSMediaPlaylist request per playlist in this account in the current Region. |
| GetMedia bandwidth | Each supported Region: 200 Megabits per second | Yes | The maximum bandwidth for GetMedia requests in megabits per second (Mbps) in this account in the current Region. |
| GetMedia concurrent connections per stream | Each supported Region: 3 | Yes | The maximum number of concurrent connections that you can make with GetMedia per stream in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetMediaForFragmentList bandwidth | Each supported Region: 200 Megabits per second | Yes | The maximum bandwidth for GetMediaForFragmentList requests in megabits per second (Mbps) in this account in the current Region. |
| GetMediaForFragmentList connections per stream | Each supported Region: 5 | No | The maximum number of concurrent connections that you can make with GetMediaForFragmentList per stream in this account in the current Region. |
| GetMediaForFragmentList fragments | Each supported Region: 1,000 | No | The maximum number of fragments per GetMediaForFragmentList request in this account in the current Region. |
| Number of Media Storage Channels | Each supported Region: 100 | Yes | The maximum number of media-storage channels that can be configured for this account in the current Region. |
| Number of signaling channels | us-east-1: 20,000<br>us-west-2: 20,000<br>Each of the other supported Regions: 10,000 | Yes | The maximum number of signaling channels that you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of video streams | us-east-1: 20,000<br><br>us-west-2: 20,000<br><br>Each of the other supported Regions: 10,000 | Yes | The maximum number of video streams that you can create in this account in the current Region. |
| PutMedia bandwidth | Each supported Region: 100 Megabits per second | Yes | The maximum bandwidth for PutMedia requests in megabits per second (Mbps) in this account in the current Region. |
| PutMedia concurrent connections per stream | Each supported Region: 1 | No | The maximum number of concurrent connections that you can make with PutMedia per stream in this account in the current Region. |
| PutMedia fragment duration | Each supported Region: 10 Seconds | Yes | The maximum fragment duration in seconds for PutMedia requests in this account in the current Region. |
| PutMedia fragment size | Each supported Region: 50 Megabytes | No | The maximum fragment size in megabytes for PutMedia requests in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| PutMedia minimum fragment duration | Each supported Region: 1 Seconds | No | The minimum fragment duration in seconds for PutMedia requests in this account in the current Region. |
| PutMedia tracks | Each supported Region: 3 | No | The maximum number of tracks within the media that is uploaded with PutMedia into this stream in this account in the current Region. |
| Rate of ConnectAsMasterAPI requests per signaling channel | Each supported Region: 3 per second | No | The maximum number of ConnectAsMaster requests that you can make per second per signaling channel in this account in the current Region. |
| Rate of ConnectAsViewerAPI requests per signaling channel | Each supported Region: 3 per second | No | The maximum number of ConnectAsViewer burst requests that you can make per second per signaling channel in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateSignalingChannelAPI requests | Each supported Region: 50 per second | Yes | The maximum number of CreateSignalingChannel requests that you can make per second in this account in the current Region. |
| Rate of CreateStreamAPI requests | Each supported Region: 50 per second | Yes | The maximum number of CreateStream requests that you can make per second in this account in the current Region. |
| Rate of DeleteEdgeConfigurationAPI requests | Each supported Region: 10 per second | No | The maximum number of DeleteEdgeConfiguration requests that you can make per second in this account in the current Region. |
| Rate of DeleteEdgeConfigurationAPI requests per resource | Each supported Region: 1 per second | No | The maximum number of DeleteEdgeConfiguration requests that you can make per second per resource in this account in the current Region. |
| Rate of DeleteSignalingChannelAPI requests | Each supported Region: 50 per second | No | The maximum number of DeleteSignalingChannel requests that you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DeleteSignalingChannelAPI requests per signaling channel | Each supported Region: 5 per second | No | The maximum number of DeleteSignalingChannel requests that you can make per second per signaling channel in this account in the current Region. |
| Rate of DeleteStreamAPI requests | Each supported Region: 50 per second | No | The maximum number of DeleteStream requests that you can make per second in this account in the current Region. |
| Rate of DeleteStreamAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of DeleteStream requests that you can make per second per stream in this account in the current Region. |
| Rate of DescribeEdgeConfigurationAPI requests | Each supported Region: 50 per second | No | The maximum number of DescribeEdgeConfiguration requests that you can make per second in this account in the current Region. |
| Rate of DescribeEdgeConfigurationAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of DescribeEdgeConfiguration requests that you can make per second per resource in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DescribeMappedResourceConfigurationAPI requests | Each supported Region: 50 per second | No | The maximum number of DescribeMappedResourceConfiguration requests that you can make per second in this account in the current Region. |
| Rate of DescribeMappedResourceConfigurationAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of DescribeMappedResourceConfiguration requests that you can make per second per resource in this account in the current Region. |
| Rate of DescribeMediaStorageConfigurationAPI requests | Each supported Region: 50 per second | No | The maximum number of DescribeMediaStorageConfiguration requests that you can make per second in this account in the current Region. |
| Rate of DescribeMediaStorageConfigurationAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of DescribeMediaStorageConfiguration requests that you can make per second per resource in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DescribeSignalingChannelAPI requests | Each supported Region: 300 per second | No | The maximum number of DescribeSignalingChannel requests that you can make per second in this account in the current Region. |
| Rate of DescribeSignalingChannelAPI requests per signaling channel | Each supported Region: 5 per second | No | The maximum number of DescribeSignalingChannel requests that you can make per second per signaling channel in this account in the current Region. |
| Rate of DescribeStreamAPI requests | Each supported Region: 300 per second | No | The maximum number of DescribeStream requests that you can make per second in this account in the current Region. |
| Rate of DescribeStreamAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of DescribeStream requests that you can make per second per stream in this account in the current Region. |
| Rate of GetDASHManifestPlaylistAPI requests per session | Each supported Region: 5 per second | No | The maximum number of GetDASHManifestPlaylist requests that you can make per second per session in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetDASHStreamingSessionURLAPI requests per stream | Each supported Region: 25 per second | No | The maximum number of GetDASHStreamingSessionURL requests that you can make per stream per second in this account in the current Region. |
| Rate of GetDataEndpointAPI requests | Each supported Region: 300 per second | No | The maximum number of GetDataEndpoint requests that you can make per second in this account in the current Region. |
| Rate of GetDataEndpointAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of GetDataEndpoint requests that you can make per second per stream in this account in the current Region. |
| Rate of GetHLSMasterPlaylistAPI requests per session | Each supported Region: 5 per second | No | The maximum number of GetHLSMasterPlaylist requests that you can make per second per session in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetHLSMediaPlaylistAPI requests per session | Each supported Region: 5 per second | No | The maximum number of GetHLSMediaPlaylist requests that you can make per second per session in this account in the current Region. |
| Rate of GetHLSStreamingSessionURLAPI requests per stream | Each supported Region: 25 per second | No | The maximum number of GetHLSStreamingSessionURL requests that you can make per stream per second in this account in the current Region. |
| Rate of GetICEServerConfigAPI requests per signaling channel | Each supported Region: 5 per second | No | The maximum number of GetICEServerConfig requests that you can make per second for this signaling channel in this account in the current Region. |
| Rate of GetMP4InitFragmentAPI requests per session | Each supported Region: 5 per second | No | The maximum number of GetMP4InitFragment requests that you can make per second per session in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetMP4MediaFragmentAPI requests per session | Each supported Region: 20 per second | No | The maximum number of GetMP4MediaFragment requests that you can make per second per session in this account in the current Region. |
| Rate of GetMediaAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of GetMedia requests that you can make per second per stream in this account in the current Region. |
| Rate of GetSignalingChannelEndpoint API requests | Each supported Region: 300 per second | No | The maximum number of GetSignalingChannelEndpoint requests that you can make per second in this account in the current Region. |
| Rate of GetSignalingChannelEndpoint API requests per signaling channel | Each supported Region: 5 per second | No | The maximum number of GetSignalingChannelEndpoint requests that you can make per second per signaling channel in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetTSFragmentAPI requests per session | Each supported Region: 20 per second | No | The maximum number of GetTSFragment requests that you can make per second per session in this account in the current Region. |
| Rate of ListEdgeAgentConfigurations API requests | Each supported Region: 50 per second | No | The maximum number of ListEdgeAgentConfigurations requests that you can make per second in this account in the current Region. |
| Rate of ListSignalingChannelsAPI requests | Each supported Region: 50 per second | No | The maximum number of ListSignalingChannels requests that you can make per second in this account in the current Region. |
| Rate of ListStreamsAPI requests | Each supported Region: 50 per second | No | The maximum number of ListStreams requests that you can make per second in this account in the current Region. |
| Rate of ListTagsForResourceAPI requests | Each supported Region: 50 per second | No | The maximum number of ListTagsForResource requests that you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListTagsForResourceAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of ListTagsForResource requests that you can make per second per resource in this account in the current Region. |
| Rate of ListTagsForStreamAPI requests | Each supported Region: 50 per second | No | The maximum number of ListTagsForStream requests that you can make per second in this account in the current Region. |
| Rate of ListTagsForStreamAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of ListTagsForStream requests that you can make per second per stream in this account in the current Region. |
| Rate of PutMediaAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of PutMedia requests that you can make per second per stream in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of SendAlexaOfferToMasterAPI requests per signaling channel | Each supported Region: 5 per second | No | The maximum number of SendAlexaOfferToMaster requests that you can make per second for this signaling channel in this account in the current Region. |
| Rate of SendICECandidateAPI requests per websocket connection | Each supported Region: 20 per second | No | The maximum number of SendICECandidate requests that you can make per second per websocket in this account in the current Region. |
| Rate of SendSDPAnswerAPI requests per websocket connection | Each supported Region: 5 per second | No | The maximum number of SendSDPAnswer requests that you can make per second per websocket in this account in the current Region. |
| Rate of SendSDPOfferAPI requests per websocket connection | Each supported Region: 5 per second | No | The maximum number of SendSDPOffer requests that you can make per second per websocket in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of StartEdgeConfigurationUpdateAPI requests | Each supported Region: 10 per second | No | The maximum number of StartEdgeConfigurationUpdate requests that you can make per second in this account in the current Region. |
| Rate of StartEdgeConfigurationUpdateAPI requests per resource | Each supported Region: 1 per second | No | The maximum number of StartEdgeConfigurationUpdate requests that you can make per second per resource in this account in the current Region. |
| Rate of TagResourceAPI requests | Each supported Region: 50 per second | No | The maximum number of TagResource requests that you can make per second in this account in the current Region. |
| Rate of TagResourceAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of TagResource requests that you can make per second per resource in this account in the current Region. |
| Rate of TagStreamAPI requests | Each supported Region: 50 per second | No | The maximum number of TagStream requests that you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of TagStreamAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of TagStream requests that you can make per second per stream in this account in the current Region. |
| Rate of UntagResourceAPI requests | Each supported Region: 50 per second | No | The maximum number of UntagResource requests that you can make per second in this account in the current Region. |
| Rate of UntagResourceAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of TagResource requests that you can make per second per resource in this account in the current Region. |
| Rate of UntagStreamAPI requests | Each supported Region: 50 per second | No | The maximum number of UntagStream requests that you can make per second in this account in the current Region. |
| Rate of UntagStreamAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of UntagStream requests that you can make per second per stream in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of UpdateDataRetentionAPI requests | Each supported Region: 50 per second | No | The maximum number of UpdateDataRetention requests that you can make per second in this account in the current Region. |
| Rate of UpdateDataRetentionAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of UpdateDataRetention requests that you can make per second per stream in this account in the current Region. |
| Rate of UpdateMediaStorageConfigurationAPI requests | Each supported Region: 10 per second | No | The maximum number of UpdateMediaStorageConfiguration requests that you can make per second in this account in the current Region. |
| Rate of UpdateMediaStorageConfigurationAPI requests per resource | Each supported Region: 5 per second | No | The maximum number of UpdateMediaStorageConfiguration requests that you can make per second per resource in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateSignalingChannelAPI requests | Each supported Region: 50 per second | No | The maximum number of UpdateSignalingChannel requests that you can make per second in this account in the current Region. |
| Rate of UpdateSignalingChannelAPI requests per signaling channel | Each supported Region: 5 per second | No | The maximum number of UpdateSignalingChannel requests that you can make per second per signaling channel in this account in the current Region. |
| Rate of UpdateStreamAPI requests | Each supported Region: 50 per second | No | The maximum number of UpdateStream requests that you can make per second in this account in the current Region. |
| Rate of UpdateStreamAPI requests per stream | Each supported Region: 5 per second | No | The maximum number of UpdateStream requests that you can make per second per stream in this account in the current Region. |
| Rate of archived fragment media per stream | Each supported Region: 500 per second | No | The maximum number of fragments that you can request media for per stream per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of archived fragment metadata per stream | Each supported Region: 10,000 per second | No | The maximum number of fragments that you can request metadata for per stream per second in this account in the current Region. |
| SendICECandidate message payload size | Each supported Region: 10 Kilobytes | No | The maximum size (in kilobytes) of SendICECandidate message payload. |
| SendSDPAnswer message payload size | Each supported Region: 10 Kilobytes | No | The maximum size (in kilobytes) of SendSDPAnswer message payload. |
| SendSDPOffer message payload size | Each supported Region: 10 Kilobytes | No | The maximum size (in kilobytes) of SendSDPOffer message payload. |
| TURN session bandwidth | Each supported Region: 5 Megabits per second | No | The maximum bandwidth in megabits per second (Mbps) that is supported per TURN session in this account in the current Region. |
| TURN session concurrent allocations per signaling channel | Each supported Region: 50 | No | The maximum number of concurrent allocated TURN sessions per signaling channel in this account in the current Region. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| TURN session expiration | Each supported Region: 300 Seconds | No | The maximum valid duration of TURN session credentials in this account in the current Region. |

For more information, see Kinesis Video Streams quotas in the *Amazon Kinesis Video Streams Developer Guide*.

# AWS Lake Formation endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | lakeformation.us-east-2.amazonaws.com | HTTPS | |
| | | lakeformation-fips.us-east-2.api.aws | HTTPS | |
| | | lakeformation-fips.us-east-2.amazonaws.com | HTTPS | |
| | | lakeformation.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | lakeformation.us-east-1.amazonaws.com | HTTPS |
| | | lakeformation-fips.us-east-1.api.aws | HTTPS |
| | | lakeformation-fips.us-east-1.amazonaws.com | HTTPS |
| | | lakeformation.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | lakeformation.us-west-1.amazonaws.com | HTTPS |
| | | lakeformation-fips.us-west-1.api.aws | HTTPS |
| | | lakeformation-fips.us-west-1.amazonaws.com | HTTPS |
| | | lakeformation.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | lakeformation.us-west-2.amazonaws.com | HTTPS |
| | | lakeformation-fips.us-west-2.api.aws | HTTPS |
| | | lakeformation-fips.us-west-2.amazonaws.com | HTTPS |
| | | lakeformation.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | lakeformation.af-south-1.amazonaws.com | HTTPS |
| | | lakeformation.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | lakeformation.ap-east-1.amazonaws.com | HTTPS |
| | | lakeformation.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | lakeformation.ap-south-2.amazonaws.com | HTTPS |
| | | lakeformation.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | lakeformation.ap-southeast-3.amazonaws.com<br>lakeformation.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | lakeformation.ap-southeast-5.amazonaws.com<br>lakeformation.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | lakeformation.ap-southeast-4.amazonaws.com<br>lakeformation.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | lakeformation.ap-south-1.amazonaws.com<br>lakeformation.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | lakeformation.ap-southeast-6.amazonaws.com<br>lakeformation.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | lakeformation.ap-northeast-3.amazonaws.com<br>lakeformation.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | lakeformation.ap-northeast-2.amazonaws.com<br>lakeformation.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | lakeformation.ap-southeast-1.amazonaws.com<br>lakeformation.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | lakeformation.ap-southeast-2.amazonaws.com<br>lakeformation.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | lakeformation.ap-east-2.amazonaws.com<br>lakeformation.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | lakeformation.ap-southeast-7.amazonaws.com<br>lakeformation.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | lakeformation.ap-northeast-1.amazonaws.com<br>lakeformation.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | lakeformation.ca-central-1.amazonaws.com<br>lakeformation-fips.ca-central-1.api.aws<br>lakeformation-fips.ca-central-1.amazonaws.com<br>lakeformation.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | lakeformation.ca-west-1.amazonaws.com | HTTPS | |
| | | lakeformation-fips.ca-west-1.api.aws | HTTPS | |
| | | lakeformation-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | lakeformation.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | lakeformation.eu-central-1.amazonaws.com | HTTPS | |
| | | lakeformation.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | lakeformation.eu-west-1.amazonaws.com | HTTPS | |
| | | lakeformation.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | lakeformation.eu-west-2.amazonaws.com | HTTPS | |
| | | lakeformation.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | lakeformation.eu-south-1.amazonaws.com | HTTPS | |
| | | lakeformation.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | lakeformation.eu-west-3.amazonaws.com | HTTPS | |
| | | lakeformation.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | lakeformation.eu-south-2.amazonaws.com | HTTPS | |
| | | lakeformation.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | lakeformation.eu-north-1.amazonaws.com | HTTPS | |
| | | lakeformation.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | lakeformation.eu-central-2.amazonaws.com | HTTPS | |
| | | lakeformation.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | lakeformation.il-central-1.amazonaws.com | HTTPS |
| | | lakeformation.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | lakeformation.mx-central-1.amazonaws.com | HTTPS |
| | | lakeformation.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | lakeformation.me-south-1.amazonaws.com | HTTPS |
| | | lakeformation.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | lakeformation.me-central-1.amazonaws.com | HTTPS |
| | | lakeformation.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | lakeformation.sa-east-1.amazonaws.com | HTTPS |
| | | lakeformation.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | lakeformation.us-gov-east-1.amazonaws.com | HTTPS |
| | | lakeformation-fips.us-gov-east-1.api.aws | HTTPS |
| | | lakeformation-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | lakeformation.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | lakeformation.us-gov-west-1.amazonaws.com | HTTPS |
| | | lakeformation-fips.us-gov-west-1.api.aws | HTTPS |
| | | lakeformation-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | lakeformation.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Length of a path that can be registered | Each supported Region: 700 | Yes | The maximum length of a path that can be registered per catalog. |
| Number of data lake administrators | Each supported Region: 30 | Yes | The maximum number of data lake administrators per catalog. |
| Number of lf tag per account | Each supported Region: 1,000 | Yes | The maximum number of lf tag per account |
| Number of lf tag policy per principal per resource type | Each supported Region: 50 | Yes | The maximum number of lf tag policy per principal per resource type |
| Number of registered paths | Each supported Region: 10,000 | Yes | The maximum number of registered paths per catalog. |
| Number of subfolders in an Amazon S3 path | Each supported Region: 20 | Yes | The maximum number of subfolders in an Amazon S3 path per catalog. |
| Number of tag values per lf tag | Each supported Region: 1,000 | Yes | The maximum number of tag values per lf tag. |

# AWS Lambda endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | lambda.us-east-2.amazonaws.com | HTTPS |
| | | lambda-fips.us-east-2.amazonaws.com | HTTPS |
| | | lambda.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | lambda.us-east-1.amazonaws.com | HTTPS |
| | | lambda-fips.us-east-1.amazonaws.com | HTTPS |
| | | lambda.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | lambda.us-west-1.amazonaws.com | HTTPS |
| | | lambda-fips.us-west-1.amazonaws.com | HTTPS |
| | | lambda.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | lambda.us-west-2.amazonaws.com | HTTPS |
| | | lambda-fips.us-west-2.amazonaws.com | HTTPS |
| | | lambda.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | lambda.af-south-1.amazonaws.com | HTTPS |
| | | lambda.af-south-1.api.aws | HTTPS |
| Asia Pacific | ap-east-1 | lambda.ap-east-1.amazonaws.com | HTTPS |
| | | lambda.ap-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | lambda.ap-south-2.amazonaws.com<br>lambda.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | lambda.ap-southeast-3.amazonaws.com<br>lambda.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | lambda.ap-southeast-5.amazonaws.com<br>lambda.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | lambda.ap-southeast-4.amazonaws.com<br>lambda.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | lambda.ap-south-1.amazonaws.com<br>lambda.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | lambda.ap-southeast-6.amazonaws.com<br>lambda.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | lambda.ap-northeast-3.amazonaws.com<br>lambda.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | lambda.ap-northeast-2.amazonaws.com<br>lambda.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | lambda.ap-southeast-1.amazonaws.com<br>lambda.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | lambda.ap-southeast-2.amazonaws.com<br>lambda.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | lambda.ap-east-2.amazonaws.com<br>lambda.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | lambda.ap-southeast-7.amazonaws.com<br>lambda.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | lambda.ap-northeast-1.amazonaws.com<br>lambda.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | lambda.ca-central-1.amazonaws.com<br>lambda.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | lambda.ca-west-1.amazonaws.com<br>lambda.ca-west-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | lambda.eu-central-1.amazonaws.com<br>lambda.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | lambda.eu-west-1.amazonaws.com<br>lambda.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | lambda.eu-west-2.amazonaws.com<br>lambda.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | lambda.eu-south-1.amazonaws.com<br>lambda.eu-south-1.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | lambda.eu-west-3.amazonaws.com<br>lambda.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | lambda.eu-south-2.amazonaws.com<br>lambda.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | lambda.eu-north-1.amazonaws.com<br>lambda.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | lambda.eu-central-2.amazonaws.com<br>lambda.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | lambda.il-central-1.amazonaws.com<br>lambda.il-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | lambda.mx-central-1.amazonaws.com | HTTPS |
| | | lambda.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | lambda.me-south-1.amazonaws.com | HTTPS |
| | | lambda.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | lambda.me-central-1.amazonaws.com | HTTPS |
| | | lambda.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | lambda.sa-east-1.amazonaws.com | HTTPS |
| | | lambda.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | lambda.us-gov-east-1.amazonaws.com | HTTPS |
| | | lambda-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | lambda.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | lambda.us-gov-west-1.amazonaws.com | HTTPS |
| | | lambda-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | lambda.us-gov-west-1.api.aws | HTTPS |

## Service quotas

**Important**

New AWS accounts have reduced concurrency and memory quotas. AWS raises these quotas automatically based on your usage. You can also request a quota increase.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Asynchronous invocation request throughput on Lambda Managed Instances | Each supported Region: 5 Megabytes/ Second | Yes | The maximum invoke request throughput (MB/s) for asynchronous invocations on Lambda Managed Instances. Asynchronous invocations occur when Lambda functions are invoked by services such as Amazon S3 or Amazon SNS. |
| Asynchronous payload | Each supported Region: 1,024 Kilobytes | No | The maximum size of an incoming asynchronous invocation request. |
| Burst rate of CreateMicrovmAuthToken API requests | Each supported Region: 50 | Yes | The maximum number of CreateMicrovmAuthToken API requests that can be sent in one burst. |
| Burst rate of CreateMicrovmShell AuthToken API requests | Each supported Region: 5 | Yes | The maximum number of CreateMicrovmShell AuthToken API requests that can be sent in one burst. |
| Burst rate of GetMicrovm API requests | Each supported Region: 100 | Yes | The maximum number of GetMicrovm API requests that can be sent in one burst. |
| Burst rate of ResumeMicrovm API requests | Each supported Region: 5 | Yes | The maximum number of ResumeMicrovm API |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | requests that can be sent in one burst. |
| Burst rate of RunMicrovm API requests | Each supported Region: 5 | Yes | The maximum number of RunMicrovm API requests that can be sent in one burst. |
| Burst rate of SuspendMicrovm API requests | Each supported Region: 2 | Yes | The maximum number of SuspendMicrovm API requests that can be sent in one burst. |
| Burst rate of TerminateMicrovm API requests | Each supported Region: 10 | Yes | The maximum number of TerminateMicrovm API requests that can be sent in one burst. |
| Capacity providers | Each supported Region: 1,000 | Yes | The maximum number of capacity providers that can be created in an account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrency scaling rate | Each supported Region: 1,000 | No | The maximum immediate increase in function concurrency that can occur in response to an increase in traffic or execution duration. Every 10 seconds, each synchronously invoked Lambda function can scale by an additional 1,000 concurrent executions until the total concurrency across your functions in a Region reaches the account concurrency limit. |
| Concurrent connections per 1 vCPU MicroVM | Each supported Region: 8 Count | No | Maximum number of concurrent connections per MicroVM with 1 vCPUs. |
| Concurrent connections per 16 vCPU MicroVM | Each supported Region: 128 Count | No | Maximum number of concurrent connections per MicroVM with 16 vCPUs. |
| Concurrent connections per 2 vCPU MicroVM | Each supported Region: 16 Count | No | Maximum number of concurrent connections per MicroVM with 2 vCPUs. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrent connections per 4 vCPU MicroVM | Each supported Region: 32 Count | No | Maximum number of concurrent connections per MicroVM with 4 vCPUs. |
| Concurrent connections per 8 vCPU MicroVM | Each supported Region: 64 Count | No | Maximum number of concurrent connections per MicroVM with 8 vCPUs. |
| Concurrent executions | Each supported Region: 1,000 Count | Yes | The maximum number of events that functions can process simultaneously in the current Region. |
| Deployment package size (console editor) | Each supported Region: 3 Megabytes | No | The maximum size of a deployment package or layer archive when you upload it through the console editor. Upload larger files with Amazon S3. |
| Deployment package size (direct upload) | Each supported Region: 50 Megabytes | No | The maximum size of a deployment package or layer archive when you upload it directly to Lambda. Upload larger files with Amazon S3. |
| Deployment package size (unzipped) | Each supported Region: 250 Megabytes | No | The maximum size of the contents of a deployment package or layer archive when its unzipped. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Durable execution storage written in megabytes | Each supported Region: 100 Megabytes | No | The maximum cumulative amount of data persisted per durable execution, including input and output payloads, checkpoints and error data, measured in megabytes. |
| DynamoDB Event Source Mapping throughput on Lambda Managed Instances | Each supported Region: 10 Megabytes/ Second | Yes | The maximum data throughput (MB/s) per DynamoDB event source mapping on Lambda Managed Instances. |
| Elastic network interfaces per VPC | Each supported Region: 3,000 | Yes | The maximum number of network interfaces that Lambda creates for a VPC with functions attached. Lambda creates a network interface for each combination of subnet and security group that functions connect to. |
| Environment variable size | Each supported Region: 4 Kilobytes | No | The maximum combined size of environment variables that are configured on a function. |
| File descriptors | Each supported Region: 1,024 | No | The maximum number of file descriptors that a function can have open. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Function and layer storage | Each supported Region: 300 Gigabytes | Yes | The amount of storage thats available for deployment packages and layer archives in the current Region. |
| Function invocation rate to initiate durable executions | Each supported Region: 300 | Yes | The maximum number of new durable executions that you can initiate per second. |
| Function layers | Each supported Region: 5 | No | The maximum number of layers that you can add to your function. |
| Function resource-based policy | Each supported Region: 20 Kilobytes | No | The maximum combined size of resource-based policies that are configured on a function. |
| Function timeout | Each supported Region: 900 | No | The maximum timeout that you can configure for a function. |
| Function versions per capacity provider | Each supported Region: 100 | No | The maximum number of function versions per capacity provider. |
| Kafka Event Source Mapping throughput in default mode on Lambda Managed Instances | Each supported Region: 1 Megabytes/Second | No | The maximum data throughput (MB/s) per Kafka event source mapping in default mode on Lambda Managed Instances. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Kafka Event Source Mappings in default mode on Lambda Managed Instances | Each supported Region: 100 | No | The maximum number of Kafka event source mappings in default mode on Lambda Managed Instances. |
| Kinesis Event Source Mapping throughput on Lambda Managed Instances | Each supported Region: 25 Megabytes/ Second | Yes | The maximum data throughput (MB/s) per Kinesis event source mapping on Lambda Managed Instances. |
| Max Execution Duration of a MicroVM (in Hours) | Each supported Region: 8 | No | The maximum execution duration of a MicroVM (in Hours). |
| Max allocated MicroVM memory | us-east-1: 1,024 Gigabytes<br>us-east-2: 1,024 Gigabytes<br>us-west-2: 1,024 Gigabytes<br>ap-northeast-1: 1,024 Gigabytes<br>Each of the other supported Regions: 400 Gigabytes | Yes | The maximum total configured baseline memory across all concurrently running MicroVMs per account per region. Individual MicroVMs can burst up to 4x their configured baseline memory. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of durable operations per durable execution | Each supported Region: 3,000 | No | The maximum number of durable operations (including customer-initiated operations and automatic retries) that trigger state persistence within a single durable execution. |
| Maximum running durable executions | us-east-1: 10,000,000<br><br>us-west-2: 10,000,000<br><br>eu-west-1: 10,000,000<br><br>Each of the other supported Regions: 5,000,000 | Yes | The maximum number of durable executions that can be running simultaneously in the current region. |
| Network bandwidth per execution environment | Each supported Region: 625 Megabits/Second | Yes | The maximum sustained network bandwidth, in megabits per second, for a Lambda functions execution environment that is not attached to a VPC. |
| Network connectors | Each supported Region: 1,000 | No | Number of Network Connectors per account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of MicroVM images | Each supported Region: 100 Count | Yes | Number of MicroVM images per account per region |
| Number of concurrent MicroVM image builds | us-east-1: 10 Count<br><br>us-east-2: 10 Count<br><br>us-west-2: 10 Count<br><br>ap-northeast-1: 10 Count<br><br>Each of the other supported Regions: 5 Count | Yes | Number of concurrent MicroVM image builds per account per region |
| Processes and threads | Each supported Region: 1,024 | No | The maximum combined number of processes and threads that a function can have open. |
| RPS per 16 vCPU MicroVm | Each supported Region: 160 Count | No | Maximum number of requests per second per MicroVM with 16 vCPUs / 32 GB. |
| RPS per 4 vCPU MicroVm | Each supported Region: 40 Count | No | Maximum number of requests per second per MicroVM with 4 vCPUs / 8 GB. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CheckpointDurableExecution API requests | us-east-1: 20,000<br>us-west-2: 20,000<br>eu-west-1: 20,000<br>Each of the other supported Regions: 10,000 | Yes | The maximum number of CheckpointDurableExecution API requests per second. |
| Rate of CreateMicrovmAuthToken API requests | Each supported Region: 50 | Yes | The maximum number of CreateMicrovmAuthToken API requests per second. |
| Rate of CreateMicrovmShellAuthToken API requests | Each supported Region: 5 | Yes | The maximum number of CreateMicrovmShellAuthToken API requests per second in this account in the current Region. |
| Rate of GetDurableExecution API requests | Each supported Region: 30 | Yes | The maximum number of GetDurableExecution API requests per second. |
| Rate of GetDurableExecutionHistory API requests | Each supported Region: 15 | Yes | The maximum number of GetDurableExecutionHistory API requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetDurableExecutionState API requests | us-east-1: 20,000<br>us-west-2: 20,000<br>eu-west-1: 20,000<br>Each of the other supported Regions: 10,000 | Yes | The maximum number of GetDurableExecutionState API requests per second. |
| Rate of GetFunction API requests | Each supported Region: 100 | No | The maximum number of GetFunction API requests per second. |
| Rate of GetMicrovm API requests | Each supported Region: 100 | Yes | The maximum number of GetMicrovm API requests per second. |
| Rate of GetPolicy API requests | Each supported Region: 15 | No | The maximum number of GetPolicy API requests per second. |
| Rate of ListDurableExecutionsByFunction API requests | Each supported Region: 15 | Yes | The maximum number of ListDurableExecutionsByFunction API requests per second. |
| Rate of ResumeMicrovm API requests | Each supported Region: 5 | Yes | The maximum number of ResumeMicrovm API requests per second. |
| Rate of RunMicrovm API requests | Each supported Region: 5 | Yes | The maximum number of RunMicrovm API requests per second |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SendDurableExecutionCallbackFailure API requests | us-east-1: 6,000<br>us-west-2: 6,000<br>eu-west-1: 6,000<br>Each of the other supported Regions: 3,000 | Yes | The maximum number of SendDurableExecutionCallbackFailure API requests per second. |
| Rate of SendDurableExecutionCallbackHeartbeat API requests | us-east-1: 6,000<br>us-west-2: 6,000<br>eu-west-1: 6,000<br>Each of the other supported Regions: 3,000 | Yes | The maximum number of SendDurableExecutionCallbackHeartbeat API requests per second. |
| Rate of SendDurableExecutionCallbackSuccess API requests | us-east-1: 6,000<br>us-west-2: 6,000<br>eu-west-1: 6,000<br>Each of the other supported Regions: 3,000 | Yes | The maximum number of SendDurableExecutionCallbackSuccess API requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StopDurableExecution API requests | us-east-1: 600<br>us-west-2: 600<br>eu-west-1: 600<br>Each of the other supported Regions: 300 | Yes | The maximum number of StopDurableExecution API requests per second. |
| Rate of SuspendMicrovm API requests | Each supported Region: 2 | Yes | The maximum number of SuspendMicrovm API requests per second. |
| Rate of TerminateMicrovm API requests | Each supported Region: 10 | Yes | The maximum number of TerminateMicrovm API requests per second. |
| Rate of capacity provider read API requests | Each supported Region: 15 | No | The maximum combined rate (requests per second) for all capacity provider read APIs. |
| Rate of capacity provider write API requests | Each supported Region: 1 | No | The maximum combined rate (requests per second) for all capacity provider write APIs. |
| Rate of control plane API requests (excludes invocation, GetFunction, and GetPolicy requests) | Each supported Region: 15 | No | The maximum number of API requests per second (excluding invocation, GetFunction, and GetPolicy requests). |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| SQS Event Source Mapping throughput in default mode on Lambda Managed Instances | Each supported Region: 5 Megabytes/ Second | No | The maximum data throughput (MB/s) per SQS event source mapping in default mode on Lambda Managed Instances. |
| Synchronous payload | Each supported Region: 6 Megabytes | No | The maximum size of an incoming synchronous invocation request or outgoing response. |
| Test events (console editor) | Each supported Region: 10 | No | The maximum amount of test events for a function through the console editor. |
| Versions per MicroVM Image | Each supported Region: 50 Count | Yes | The maximum number of versions per MicroVM image per account per region. |

For more information, see Lambda quotas in the *AWS Lambda Developer Guide*.

## AWS Launch Wizard endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | launchwizard.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | launchwizard.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | launchwizard.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | launchwizard.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | launchwizard.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | launchwizard.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | launchwizard.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | launchwizard.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | launchwizard.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | launchwizard.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | launchwizard.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | launchwizard.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | launchwizard.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | launchwizard.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | launchwizard.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | launchwizard.ap-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | launchwizard.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | launchwizard.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | launchwizard.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | launchwizard.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | launchwizard.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | launchwizard.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | launchwizard.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | launchwizard.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | launchwizard.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | launchwizard.eu-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | launchwizard.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | launchwizard.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | launchwizard.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | launchwizard.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | launchwizard.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | launchwizard.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | launchwizard.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | launchwizard.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | launchwizard.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Active deployments | Each supported Region: 75 | Yes | The maximum number of deployments with a status other than DELETED in this account in the current Region. |
| Deployments | Each supported Region: 300 | Yes | The maximum number of deployments that you can have in this account in the current Region. |
| In-Progress Deployments | Each supported Region: 10 | Yes | The maximum number of in-progress deployments with a status of IN_PROGRESS or DELETE_IN_PROGRESS in this account in the current Region. |
| Settings Set | Each supported Region: 50 | Yes | The maximum number of settings sets that you can create in this account in the current Region. |

# Amazon Lex endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# V2 service endpoints

## Model building endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | models-v2-lex.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | models-v2-lex.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | models-v2-lex.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | models-v2-lex.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | models-v2-lex.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | models-v2-lex.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | models-v2-lex.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | models-v2-lex.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | models-v2-lex.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | models-v2-lex.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | models-v2-lex.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | models-v2-lex.us-gov-west-1.amazonaws.com | HTTPS | |

## Runtime endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | runtime-v2-lex.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | runtime-v2-lex.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | runtime-v2-lex.af-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | runtime-v2-lex.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | runtime-v2-lex.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | runtime-v2-lex.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | runtime-v2-lex.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | runtime-v2-lex.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | runtime-v2-lex.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | runtime-v2-lex.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | runtime-v2-lex.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | runtime-v2-lex.us-gov-west-1.amazonaws.com | HTTPS | |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Bot channel associations per bot alias (V2) | Each supported Region: 10 | No | The maximum number of bot channel associations that you can create per bot alias in this account in the current Region. |
| Bots per account (V2) | Each supported Region: 100 | Yes | The maximum number of bots that you can create in this account in the current Region. |
| Characters per custom slot type value (V2) | Each supported Region: 500 | No | The maximum number of characters that you can have per custom slot type value in this account in the current Region. |
| Characters per sample utterance (V2) | Each supported Region: 500 | No | The maximum number of characters that you can have per intent or slot sample utterance in this account in the current Region. |
| Custom slot type values and synonyms per bot locale (V2) | Each supported Region: 50,000 | No | The maximum number of custom slot type values and synonyms that you can have per locale per bot in this account in the current Region. |
| Custom slot types per bot locale (V2) | Each supported Region: 100 | No | The maximum number of custom slot types that |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | you can create per locale per bot in this account in the current Region. |
| Sample utterances per intent (V2) | Each supported Region: 1,500 | Yes | The maximum number of sample utterances that you can create per intent in this account in the current Region. |
| Sample utterances per slot (V2) | Each supported Region: 10 | Yes | The maximum number of sample utterances that you can create per slot in this account in the current Region. |
| Slots per bot locale (V2) | Each supported Region: 2,000 | No | The maximum number of slots that you can create per locale per bot in this account in the current Region. |
| Slots per intent (V2) | Each supported Region: 100 | No | The maximum number of slots that you can create per intent in this account in the current Region. |
| Total characters in sample utterances per bot locale (V2) | Each supported Region: 200,000 | No | The maximum number of characters that you can use per locale per bot for all intent and slot sample utterances in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Values and synonyms per custom slot type (V2) | Each supported Region: 10,000 | No | The maximum number of values and synonyms that you can have per custom slot type in this account in the current Region. |
| Versions per bot (V2) | Each supported Region: 100 | No | The maximum number of versions that you can create per bot in this account in the current Region. |

# AWS License Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Topics

- Endpoints for working with licenses
- Endpoints for working with user-based subscriptions
- Endpoints for working with Linux subscriptions

## Endpoints for working with licenses

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | license-manager.us-east-2.amazonaws.com<br>license-manager-fips.us-east-2.amazonaws.com<br>license-manager.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | license-manager.us-east-1.amazonaws.com<br>license-manager-fips.us-east-1.amazonaws.com<br>license-manager.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | license-manager.us-west-1.amazonaws.com<br>license-manager-fips.us-west-1.amazonaws.com<br>license-manager.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | license-manager.us-west-2.amazonaws.com<br>license-manager-fips.us-west-2.amazonaws.com<br>license-manager.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | license-manager.af-south-1.amazonaws.com<br>license-manager.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific | ap-east-1 | license-manager.ap-east-1.amazonaws.com<br>license-manager.ap-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | license-manager.ap-south-2.amazonaws.com<br>license-manager.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | license-manager.ap-southeast-3.amazonaws.com<br>license-manager.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | license-manager.ap-southeast-5.amazonaws.com<br>license-manager.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | license-manager.ap-southeast-4.amazonaws.com<br>license-manager.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | license-manager.ap-south-1.amazonaws.com<br>license-manager.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | license-manager.ap-southeast-6.amazonaws.com<br>license-manager.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | license-manager.ap-northeast-3.amazonaws.com<br>license-manager.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | license-manager.ap-northeast-2.amazonaws.com<br>license-manager.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | license-manager.ap-southeast-1.amazonaws.com<br>license-manager.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | license-manager.ap-southeast-2.amazonaws.com<br>license-manager.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | license-manager.ap-east-2.amazonaws.com<br>license-manager.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | license-manager.ap-southeast-7.amazonaws.com<br>license-manager.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | license-manager.ap-northeast-1.amazonaws.com<br>license-manager.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | license-manager.ca-central-1.amazonaws.com<br>license-manager.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | license-manager.ca-west-1.amazonaws.com<br>license-manager.ca-west-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | license-manager.eu-central-1.amazonaws.com | HTTPS |
| | | license-manager.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | license-manager.eu-west-1.amazonaws.com | HTTPS |
| | | license-manager.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | license-manager.eu-west-2.amazonaws.com | HTTPS |
| | | license-manager.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | license-manager.eu-south-1.amazonaws.com | HTTPS |
| | | license-manager.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | license-manager.eu-west-3.amazonaws.com | HTTPS |
| | | license-manager.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | license-manager.eu-south-2.amazonaws.com | HTTPS |
| | | license-manager.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | license-manager.eu-north-1.amazonaws.com | HTTPS |
| | | license-manager.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | license-manager.eu-central-2.amazonaws.com | HTTPS |
| | | license-manager.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | license-manager.il-central-1.amazonaws.com | HTTPS |
| | | license-manager.il-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | license-manager.mx-central-1.amazonaws.com<br>license-manager.mx-central-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | license-manager.me-south-1.amazonaws.com<br>license-manager.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | license-manager.me-central-1.amazonaws.com<br>license-manager.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | license-manager.sa-east-1.amazonaws.com<br>license-manager.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | license-manager.us-gov-east-1.amazonaws.com<br>license-manager-fips.us-gov-east-1.amazonaws.com<br>license-manager.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | license-manager.us-gov-west-1.amazonaws.com<br>license-manager-fips.us-gov-west-1.amazonaws.com<br>license-manager.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |

## Endpoints for working with user-based subscriptions

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | license-manager-user-subscriptions.us-east-2.amazonaws.com<br>license-manager-user-subscriptions-fips.us-east-2.amazonaws.com<br>license-manager-user-subscriptions.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | license-manager-user-subscriptions.us-east-1.amazonaws.com<br>license-manager-user-subscriptions-fips.us-east-1.amazonaws.com<br>license-manager-user-subscriptions.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | license-manager-user-subscriptions.us-west-1.amazonaws.com<br>license-manager-user-subscriptions-fips.us-west-1.amazonaws.com<br>license-manager-user-subscriptions.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | license-manager-user-subscriptions.us-west-2.amazonaws.com<br>license-manager-user-subscriptions-fips.us-west-2.amazonaws.com<br>license-manager-user-subscriptions.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | license-manager-user-subscriptions.af-south-1.amazonaws.com<br><br>license-manager-user-subscriptions.af-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | license-manager-user-subscriptions.ap-east-1.amazonaws.com<br><br>license-manager-user-subscriptions.ap-east-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | license-manager-user-subscriptions.ap-south-2.amazonaws.com<br><br>license-manager-user-subscriptions.ap-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | license-manager-user-subscriptions.ap-southeast-3.amazonaws.com<br><br>license-manager-user-subscriptions.ap-southeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | license-manager-user-subscriptions.ap-southeast-5.amazonaws.com<br><br>license-manager-user-subscriptions.ap-southeast-5.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | license-manager-user-subscriptions.ap-southeast-4.amazonaws.com<br><br>license-manager-user-subscriptions.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | license-manager-user-subscriptions.ap-south-1.amazonaws.com<br><br>license-manager-user-subscriptions.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | license-manager-user-subscriptions.ap-northeast-3.amazonaws.com<br><br>license-manager-user-subscriptions.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | license-manager-user-subscriptions.ap-northeast-2.amazonaws.com<br><br>license-manager-user-subscriptions.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | license-manager-user-subscriptions.ap-southeast-1.amazonaws.com<br><br>license-manager-user-subscriptions.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | license-manager-user-subscriptions.ap-southeast-2.amazonaws.com<br><br>license-manager-user-subscriptions.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | license-manager-user-subscriptions.ap-east-2.amazonaws.com<br><br>license-manager-user-subscriptions.ap-east-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | license-manager-user-subscriptions.ap-southeast-7.amazonaws.com<br><br>license-manager-user-subscriptions.ap-southeast-7.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | license-manager-user-subscriptions.ap-northeast-1.amazonaws.com<br><br>license-manager-user-subscriptions.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | license-manager-user-subscriptions.ca-central-1.amazonaws.com<br><br>license-manager-user-subscriptions.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | license-manager-user-subscriptions.ca-west-1.amazonaws.com<br><br>license-manager-user-subscriptions.ca-west-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | license-manager-user-subscriptions.eu-central-1.amazonaws.com<br><br>license-manager-user-subscriptions.eu-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | license-manager-user-subscriptions.eu-west-1.amazonaws.com<br><br>license-manager-user-subscriptions.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | license-manager-user-subscriptions.eu-west-2.amazonaws.com<br><br>license-manager-user-subscriptions.eu-west-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Milan) | eu-south-1 | license-manager-user-subscriptions.eu-south-1.amazonaws.com<br><br>license-manager-user-subscriptions.eu-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Paris) | eu-west-3 | license-manager-user-subscriptions.eu-west-3.amazonaws.com<br><br>license-manager-user-subscriptions.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Spain) | eu-south-2 | license-manager-user-subscriptions.eu-south-2.amazonaws.com<br><br>license-manager-user-subscriptions.eu-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | license-manager-user-subscriptions.eu-north-1.amazonaws.com<br><br>license-manager-user-subscriptions.eu-north-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Zurich) | eu-central-2 | license-manager-user-subscriptions.eu-central-2.amazonaws.com<br><br>license-manager-user-subscriptions.eu-central-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | license-manager-user-subscriptions.il-central-1.amazonaws.com<br><br>license-manager-user-subscriptions.il-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Mexico (Central) | mx-central-1 | license-manager-user-subscriptions.mx-central-1.amazonaws.com<br><br>license-manager-user-subscriptions.mx-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | license-manager-user-subscriptions.me-south-1.amazonaws.com<br><br>license-manager-user-subscriptions.me-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (UAE) | me-central-1 | license-manager-user-subscriptions.me-central-1.amazonaws.com<br><br>license-manager-user-subscriptions.me-central-1.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | license-manager-user-subscriptions.sa-east-1.amazonaws.com<br><br>license-manager-user-subscriptions.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | license-manager-user-subscriptions.us-gov-east-1.amazonaws.com<br><br>license-manager-user-subscriptions.us-gov-east-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | license-manager-user-subscriptions.us-gov-west-1.amazonaws.com<br><br>license-manager-user-subscriptions.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

## Endpoints for working with Linux subscriptions

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | license-manager-linux-subscriptions.us-east-2.amazonaws.com<br><br>license-manager-linux-subscriptions-fips.us-east-2.amazonaws.com<br><br>license-manager-linux-subscriptions.us-east-2.api.aws | HTTPS<br><br>HTTPS<br><br>HTTPS | |
| US East (N. Virginia) | us-east-1 | license-manager-linux-subscriptions.us-east-1.amazonaws.com<br><br>license-manager-linux-subscriptions-fips.us-east-1.amazonaws.com<br><br>license-manager-linux-subscriptions.us-east-1.api.aws | HTTPS<br><br>HTTPS<br><br>HTTPS | |
| US West (N. California) | us-west-1 | license-manager-linux-subscriptions.us-west-1.amazonaws.com<br><br>license-manager-linux-subscriptions-fips.us-west-1.amazonaws.com | HTTPS<br><br>HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | license-manager-linux-subscriptions.us-west-1.api.aws | | |
| US West (Oregon) | us-west-2 | license-manager-linux-subscriptions.us-west-2.amazonaws.com<br>license-manager-linux-subscriptions-fips.us-west-2.amazonaws.com<br>license-manager-linux-subscriptions.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | license-manager-linux-subscriptions.af-south-1.amazonaws.com<br>license-manager-linux-subscriptions.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | license-manager-linux-subscriptions.ap-east-1.amazonaws.com<br>license-manager-linux-subscriptions.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | license-manager-linux-subscriptions.ap-south-2.amazonaws.com<br>license-manager-linux-subscriptions.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | license-manager-linux-subscriptions.ap-southeast-3.amazonaws.com<br>license-manager-linux-subscriptions.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | license-manager-linux-subscriptions.ap-southeast-5.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-southeast-5.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | license-manager-linux-subscriptions.ap-southeast-4.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | license-manager-linux-subscriptions.ap-south-1.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | license-manager-linux-subscriptions.ap-southeast-6.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-southeast-6.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | license-manager-linux-subscriptions.ap-northeast-3.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | license-manager-linux-subscriptions.ap-northeast-2.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | license-manager-linux-subscriptions.ap-southeast-1.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | license-manager-linux-subscriptions.ap-southeast-2.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | license-manager-linux-subscriptions.ap-east-2.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-east-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | license-manager-linux-subscriptions.ap-southeast-7.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-southeast-7.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | license-manager-linux-subscriptions.ap-northeast-1.amazonaws.com<br><br>license-manager-linux-subscriptions.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | license-manager-linux-subscriptions.ca-central-1.amazonaws.com<br><br>license-manager-linux-subscriptions.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | license-manager-linux-subscriptions.ca-west-1.amazonaws.com<br>license-manager-linux-subscriptions.ca-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | license-manager-linux-subscriptions.eu-central-1.amazonaws.com<br>license-manager-linux-subscriptions.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | license-manager-linux-subscriptions.eu-west-1.amazonaws.com<br>license-manager-linux-subscriptions.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | license-manager-linux-subscriptions.eu-west-2.amazonaws.com<br>license-manager-linux-subscriptions.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | license-manager-linux-subscriptions.eu-south-1.amazonaws.com<br>license-manager-linux-subscriptions.eu-south-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | license-manager-linux-subscriptions.eu-west-3.amazonaws.com<br>license-manager-linux-subscriptions.eu-west-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | license-manager-linux-subscriptions.eu-south-2.amazonaws.com<br><br>license-manager-linux-subscriptions.eu-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | license-manager-linux-subscriptions.eu-north-1.amazonaws.com<br><br>license-manager-linux-subscriptions.eu-north-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Zurich) | eu-central-2 | license-manager-linux-subscriptions.eu-central-2.amazonaws.com<br><br>license-manager-linux-subscriptions.eu-central-2.api.aws | HTTPS<br><br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | license-manager-linux-subscriptions.il-central-1.amazonaws.com<br><br>license-manager-linux-subscriptions.il-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Mexico (Central) | mx-central-1 | license-manager-linux-subscriptions.mx-central-1.amazonaws.com<br><br>license-manager-linux-subscriptions.mx-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | license-manager-linux-subscriptions.me-south-1.amazonaws.com<br><br>license-manager-linux-subscriptions.me-south-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | license-manager-linux-subscriptions.me-central-1.amazonaws.com<br><br>license-manager-linux-subscriptions.me-central-1.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | license-manager-linux-subscriptions.sa-east-1.amazonaws.com<br><br>license-manager-linux-subscriptions.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | license-manager-linux-subscriptions.us-gov-east-1.amazonaws.com<br><br>license-manager-linux-subscriptions.us-gov-east-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | license-manager-linux-subscriptions.us-gov-west-1.amazonaws.com<br><br>license-manager-linux-subscriptions.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

### Topics

- Quotas for working with licenses
- Quotas for working with user-based subscriptions
- Quotas for working with Linux subscriptions

## Quotas for working with licenses

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Custom license asset rulesets per account | Each supported Region: 25 | No | Total number of custom license asset rulesets that can be created per account. |
| Extend license consumption per consumption token | Each supported Region: 1 | No | The number of ExtendLicenseConsumption calls that can be made for a consumption token per hour. |
| GetAccessTokens calls | Each supported Region: 10 | No | The total number of GetAccessToken calls that can be made for a license per hour. |
| License asset groups per account | Each supported Region: 25 | No | Total number of license asset groups that can be created per account. |
| License asset rulesets per asset group | Each supported Region: 5 | No | Total number of license asset rulesets that can attached to a license asset group. |
| License configuration associations per resource | Each supported Region: 10 | Yes | Number of license configurations that can be associated with a resource. |
| License configurations | Each supported Region: 25 | Yes | Total number of license configurations that can be created in an account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| License conversion tasks per resource per day | Each supported Region: 5 | Yes | Number of license conversion tasks that can be created per resource per day. |
| Maximum number of concurrent organization grant activities | Each supported Region: 10 | No | The maximum number of concurrent grant activities for an organization. |
| Number of Report generators | Each supported Region: 25 | No | The Total number of report generators that can be created. |
| Number of account discovery mode updates per day | Each supported Region: 1 | No | The maximum number of account discovery mode updates per day for an account. |
| Number of accounts per organization for License Manager | Each supported Region: 15,000 | No | The maximum number of accounts within your organization that License Manager can support. |
| Number of grants per license | Each supported Region: 2,000 | No | The total number of active grants per license. |
| Number of instances aggregated per license asset group per account | Each supported Region: 10,000 | Yes | Total number of instances each license asset group can aggregate usage data for each account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of instances aggregated per license asset group per organization | Each supported Region: 50,000 | Yes | Total number of instances each license asset group can aggregate usage data for each organization. |
| Number of instances tracked per license asset group per account per region | Each supported Region: 5,000 | Yes | Total number of instances each license asset group can track and evaluate for each account in a region. |
| Number of instances tracked per license asset group per organization per region | Each supported Region: 25,000 | Yes | Total number of instances each license asset group can track and evaluate for each organization in a region. |
| Number of licenses you can create | Each supported Region: 2,000 | No | The total number of licenses that can be created in an account. |
| Number of received licenses per product | Each supported Region: 10 | No | The total number of licenses received with same product SKU. |
| Number of tokens per account and license | Each supported Region: 10 | No | The total number of tokens per license that can be created in an account. |
| Number of updates for a report generator per day | Each supported Region: 25 | No | The maximum number of updates per day for a given report generator. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rules per custom license asset ruleset | Each supported Region: 5 | No | Total number of rules per custom license asset ruleset. |
| Total number counted entitlements per checkout | Each supported Region: 5 | No | The total number of counted entitlements that can be specified on a single CheckoutLicense call. |
| Total number counted entitlements per license | Each supported Region: 25 | No | The total number of counted entitlements allowed per license. |
| Total number uncounted entitlements per license | Each supported Region: 25 | No | The total number of uncounted entitlements allowed per license. |

## Quotas for working with user-based subscriptions

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Instance associations per user | Each supported Region: 25 | Yes | The maximum number of instances that a user can be associated with in this account in the current Region. |
| User-based subscriptions for Office Professional Plus | Each supported Region: 50 | Yes | The maximum number of user-based subscriptions for Office Professional Plus that you can |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | create in this account in the current Region. |
| User-based subscriptions for Visual Studio Enterprise | Each supported Region: 30 | Yes | The maximum number of user-based subscriptions for Visual Studio Enterprise that you can create in this account in the current Region. |
| User-based subscriptions for Visual Studio Professional | Each supported Region: 50 | Yes | The maximum number of user-based subscriptions for Visual Studio Professional that you can create in this account in the current Region. |

### Quotas for working with Linux subscriptions

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of discovered resources | Each supported Region: 5,000 | No | The total number of resources (EC2 instances) that can be discovered and displayed. |

## Amazon Lightsail endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Note**

- Lightsail for Research is available in the same AWS Regions as the Amazon Lightsail service with the exception of the US East (N. Virginia) Region. Lightsail for Research also uses the same endpoints as Lightsail.
- The Amazon Lightsail service offers dual stack endpoints so that you can send it either IPv4 or IPv6 requests. In the following table, dual stack endpoints resemble `lightsail.region.api.aws` and IPv4 endpoints resemble `lightsail.region.amazonaws.com`. For more information, see Dual stack endpoints.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | lightsail.us-east-2.amazonaws.com | HTTPS |
| | | lightsail.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | lightsail.us-east-1.amazonaws.com | HTTPS |
| | | lightsail.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | lightsail.us-west-2.amazonaws.com | HTTPS |
| | | lightsail.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | lightsail.ap-east-1.amazonaws.com | HTTPS |
| | | lightsail.ap-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | lightsail.ap-southeast-3.amazonaws.com<br>lightsail.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | lightsail.ap-southeast-5.amazonaws.com<br>lightsail.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | lightsail.ap-south-1.amazonaws.com<br>lightsail.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | lightsail.ap-northeast-2.amazonaws.com<br>lightsail.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | lightsail.ap-southeast-1.amazonaws.com<br>lightsail.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | lightsail.ap-southeast-2.amazonaws.com<br>lightsail.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | lightsail.ap-northeast-1.amazonaws.com<br>lightsail.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | lightsail.ca-central-1.amazonaws.com<br>lightsail.ca-central-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | lightsail.eu-central-1.amazonaws.com | HTTPS |
| | | lightsail.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | lightsail.eu-west-1.amazonaws.com | HTTPS |
| | | lightsail.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | lightsail.eu-west-2.amazonaws.com | HTTPS |
| | | lightsail.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | lightsail.eu-west-3.amazonaws.com | HTTPS |
| | | lightsail.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | lightsail.eu-south-2.amazonaws.com | HTTPS |
| | | lightsail.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | lightsail.eu-north-1.amazonaws.com | HTTPS |
| | | lightsail.eu-north-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | lightsail.sa-east-1.amazonaws.com | HTTPS |
| | | lightsail.sa-east-1.api.aws | HTTPS |

## Service quotas

New AWS accounts might start with quotas that are lower than those described here.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Allowed cookies per cache behavior for a distribution | Each supported Region: 10 | No | The maximum number of allowed cookies per cache behavior. |
| Allowed headers per cache behavior for a distribution | Each supported Region: 10 | No | The maximum number of allowed headers per cache behavior. |
| Allowed query strings per cache behavior for a distribution | Each supported Region: 10 | No | The maximum number of allowed query strings per cache behavior. |
| Block storage disks per instance | Each supported Region: 15 | No | The maximum number of block storage disks that can be attached per instance. |
| Container service certificates | Each supported Region: 4 | No | The maximum number of certificates that can be attached per container service. |
| Container service custom domains | Each supported Region: 4 | No | The maximum number of custom domains per container service. |
| Container service deployment containers | Each supported Region: 10 | No | The maximum number of containers per deployment on a container service. |
| Container service deployment versions | Each supported Region: 50 | No | The maximum number of deployment versions per container service. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Container service logs storage days | Each supported Region: 4 | No | The maximum number of days that container service logs are stored. |
| Container service nodes | Each supported Region: 20 | No | The maximum number of nodes per container service. |
| Container service stored container images | Each supported Region: 150 | No | The maximum number of stored container images per container service. |
| Container services | Each supported Region: 100 | No | The maximum number of container services per region. |
| Custom domain names per distribution | Each supported Region: 10 | No | The maximum number of custom domain names per distribution. You can specify up to 10 domains for a Lightsail SSL/TLS certificate. You can then use the certificate to enable custom domains on your Lightsail distribution. |
| DNS zones (or domains) | Each supported Region: 6 | No | The maximum number of DNS zones (or domains) per account. |
| Data transfer rate per distribution | Each supported Region: 150 | No | The maximum data transfer rate (in GB per second) per distribution. |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Databases | Each supported Region: 40 | No | The maximum number of databases per region. |
| Default behaviors (default cache behavior) per distribution | Each supported Region: 1 | No | The maximum number of default behaviors (default cache behavior) per distribution. |
| Directory and file overrides per distribution | Each supported Region: 25 | No | The maximum number of directory and file overrides per distribution. You can create up to 25 directory and file overrides, and you can specify one cache behavior per override. |
| Distributions | Each supported Region: 20 | No | The maximum number of distributions per account. |
| Instances | Each supported Region: 20 | Yes | The maximum number of vCPUs per Region for Lightsail instances. |
| Load balancers | Each supported Region: 10 | No | The maximum number of load balancers per region. |
| Maximum active certificates | Each supported Region: 10 | No | The maximum number of active certificates per region, per account. |
| Maximum block storage disk space | Each supported Region: 16,000 Gigabytes | No | The maximum amount of disk space (in GB) per block storage disk. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Maximum buckets per account | Each supported Region: 20 | No | The maximum number of buckets per account |
| Maximum certificates | Each supported Region: 20 | No | The maximum number of certificates per region, per account in the last 365 days. |
| Maximum keys per bucket | Each supported Region: 2 | No | The maximum number of keys per bucket |
| Minimum block storage disk space | Each supported Region: 8 Gigabytes | No | The minimum amount of disk space (in GB) per block storage disk. |
| Origins per distribution | Each supported Region: 1 | No | The maximum number of origins per distribution. You can specify one instance or one load balancer per distribution. |
| Parallel RDP connections using the browser-based RDP client | Each supported Region: 1 | No | The maximum number of parallel RDP connections using the browser-based RDP client per region, per account. |
| Parallel SSH connections using the browser-based SSH client | Each supported Region: 5 | No | The maximum number of parallel SSH connections using the browser-based SSH client per region, per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Response timeout per origin for a distribution | Each supported Region: 60 Seconds | No | The response timeout per origin (4-60 seconds). |
| Static IP addresses | Each supported Region: 5 | Yes | The maximum number of static IP addresses per region. |
| Tags | Each supported Region: 50 | No | The maximum number of tags per resource. |
| Total attached block storage disk space | Each supported Region: 20,000 Gigabytes | No | The maximum amount of attached block storage disk space (in GB) per region. |

# Amazon Location Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Amazon Location is available in a number of AWS Regions. See the sections below for details:

- the section called “IPv4 and IPv6”
- the section called “IPv4 only”

**Note**

We use the following service codes:

| Service | Service code |
| --- | --- |
| Amazon Location Maps | maps |
| Amazon Location Places | places |
| Amazon Location Routes | routes |
| Amazon Location Geofences | geofencing |
| Amazon Location Trackers | tracking |
| Amazon Location Metadata | metadata |

## IPv4 and IPv6

The general syntax for the Amazon Location regional dual-stack endpoint, which supports both IPv4 and IPv6, is as follows:

```
protocol://service-code.geo.region-code.api.aws
```

For example, the regional endpoint for Amazon Location Maps for US East (N. Virginia) is: `https://maps.geo.us-east-1.api.aws`.

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | **Maps:** maps.geo.us-east-2.api.aws<br><br>**Places:** places.geo.us-east-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| | | **Routes:** routes.geo.us-east-2.api.aws<br><br>**Geofences:** geofencing.geo.us-east-2.api.aws<br><br>**Trackers:** tracking.geo.us-east-2.api.aws<br><br>**Metadata:** metadata.geo.us-east-2.api.aws | |
| US East (N. Virginia) | us-east-1 | **Maps:** maps.geo.us-east-1.api.aws<br><br>**Places:** places.geo.us-east-1.api.aws<br><br>**Routes:** routes.geo.us-east-1.api.aws<br><br>**Geofences:** geofencing.geo.us-east-1.api.aws<br><br>**Trackers:** tracking.geo.us-east-1.api.aws<br><br>**Metadata:** metadata.geo.us-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US West (Oregon) | us-west-2 | **Maps:** maps.geo.us-west-2.api.aws<br><br>**Places:** places.geo.us-west-2.api.aws<br><br>**Routes:** routes.geo.us-west-2.api.aws<br><br>**Geofences:** geofencing.geo.us-west-2.api.aws<br><br>**Trackers:** tracking.geo.us-west-2.api.aws<br><br>**Metadata:** metadata.geo.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Malaysia) | ap-southeast-5 | **Maps:** maps.geo.ap-southeast-5.api.aws<br><br>**Places:** places.geo.ap-southeast-5.api.aws<br><br>**Routes:** routes.geo.ap-southeast-5.api.aws<br><br>**Geofences:** geofencing.geo.ap-southeast-5.api.aws<br><br>**Trackers:** tracking.geo.ap-southeast-5.api.aws<br><br>**Metadata:** metadata.geo.ap-southeast-5.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Mumbai) | ap-south-1 | **Maps:** maps.geo.ap-south-1.api.aws<br><br>**Places:** places.geo.ap-south-1.api.aws<br><br>**Routes:** routes.geo.ap-south-1.api.aws<br><br>**Geofences:** geofencing.geo.ap-south-1.api.aws<br><br>**Trackers:** tracking.geo.ap-south-1.api.aws<br><br>**Metadata:** metadata.geo.ap-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Singapore) | ap-southeast-1 | **Maps:** maps.geo.ap-southeast-1.api.aws<br><br>**Places:** places.geo.ap-southeast-1.api.aws<br><br>**Routes:** routes.geo.ap-southeast-1.api.aws<br><br>**Geofences:** geofencing.geo.ap-southeast-1.api.aws<br><br>**Trackers:** tracking.geo.ap-southeast-1.api.aws<br><br>**Metadata:** metadata.geo.ap-southeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Sydney) | ap-southeast-2 | **Maps:** maps.geo.ap-southeast-2.api.aws<br><br>**Places:** places.geo.ap-southeast-2.api.aws<br><br>**Routes:** routes.geo.ap-southeast-2.api.aws<br><br>**Geofences:** geofencing.geo.ap-southeast-2.api.aws<br><br>**Trackers:** tracking.geo.ap-southeast-2.api.aws<br><br>**Metadata:** metadata.geo.ap-southeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Tokyo) | ap-northeast-1 | **Maps:** maps.geo.ap-northeast-1.api.aws<br><br>**Places:** places.geo.ap-northeast-1.api.aws<br><br>**Routes:** routes.geo.ap-northeast-1.api.aws<br><br>**Geofences:** geofencing.geo.ap-northeast-1.api.aws<br><br>**Trackers:** tracking.geo.ap-northeast-1.api.aws<br><br>**Metadata:** metadata.geo.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Canada (Central) | ca-central-1 | **Maps:** maps.geo.ca-central-1.api.aws<br><br>**Places:** places.geo.ca-central-1.api.aws<br><br>**Routes:** routes.geo.ca-central-1.api.aws<br><br>**Geofences:** geofencing.geo.ca-central-1.api.aws<br><br>**Trackers:** tracking.geo.ca-central-1.api.aws<br><br>**Metadata:** metadata.geo.ca-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Frankfurt) | eu-central-1 | **Maps:** maps.geo.eu-central-1.api.aws<br><br>**Places:** places.geo.eu-central-1.api.aws<br><br>**Routes:** routes.geo.eu-central-1.api.aws<br><br>**Geofences:** geofencing.geo.eu-central-1.api.aws<br><br>**Trackers:** tracking.geo.eu-central-1.api.aws<br><br>**Metadata:** metadata.geo.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Ireland) | eu-west-1 | **Maps:** maps.geo.eu-west-1.api.aws<br><br>**Places:** places.geo.eu-west-1.api.aws<br><br>**Routes:** routes.geo.eu-west-1.api.aws<br><br>**Geofences:** geofencing.geo.eu-west-1.api.aws<br><br>**Trackers:** tracking.geo.eu-west-1.api.aws<br><br>**Metadata:** metadata.geo.eu-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (London) | eu-west-2 | **Maps:** maps.geo.eu-west-2.api.aws<br><br>**Places:** places.geo.eu-west-2.api.aws<br><br>**Routes:** routes.geo.eu-west-2.api.aws<br><br>**Geofences:** geofencing.geo.eu-west-2.api.aws<br><br>**Trackers:** tracking.geo.eu-west-2.api.aws<br><br>**Metadata:** metadata.geo.eu-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Spain) | eu-south-2 | **Maps:** maps.geo.eu-south-2.api.aws<br><br>**Places:** places.geo.eu-south-2.api.aws<br><br>**Routes:** routes.geo.eu-south-2.api.aws<br><br>**Geofences:** geofencing.geo.eu-south-2.api.aws<br><br>**Trackers:** tracking.geo.eu-south-2.api.aws<br><br>**Metadata:** metadata.geo.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Stockholm) | eu-north-1 | **Maps:** maps.geo.eu-north-1.api.aws<br><br>**Places:** places.geo.eu-north-1.api.aws<br><br>**Routes:** routes.geo.eu-north-1.api.aws<br><br>**Geofences:** geofencing.geo.eu-north-1.api.aws<br><br>**Trackers:** tracking.geo.eu-north-1.api.aws<br><br>**Metadata:** metadata.geo.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| South America (São Paulo) | sa-east-1 | **Maps:** maps.geo.sa-east-1.api.aws<br><br>**Places:** places.geo.sa-east-1.api.aws<br><br>**Routes:** routes.geo.sa-east-1.api.aws<br><br>**Geofences:** geofencing.geo.sa-east-1.api.aws<br><br>**Trackers:** tracking.geo.sa-east-1.api.aws<br><br>**Metadata:** metadata.geo.sa-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-West) | us-gov-west-1 | **Maps:** maps.geo.us-gov-west-1.api.aws<br><br>**Places:** places.geo.us-gov-west-1.api.aws<br><br>**Routes:** routes.geo.us-gov-west-1.api.aws<br><br>**Geofences:** geofencing.geo.us-gov-west-1.api.aws<br><br>**Trackers:** tracking.geo.us-gov-west-1.api.aws<br><br>**Metadata:** metadata.geo.us-gov-west-1.api.aws<br><br>**FIPS**<br><br>**Maps:** maps.geo-fips.us-gov-west-1.api.aws<br><br>**Places:** places.geo-fips.us-gov-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| | | **Routes:** routes.geo-fips.us-gov-west-1.api.aws<br><br>**Geofences:** geofencing.geo-fips.us-gov-west-1.api.aws<br><br>**Trackers:** tracking.geo-fips.us-gov-west-1.api.aws<br><br>**Metadata:** metadata.geo-fips.us-gov-west-1.api.aws | |

### IPv4 only

The existing regional endpoint, which supports only IPv4, is as follows:

```
protocol://service-code.geo.region-code.amazonaws.com
```

For example, the regional endpoint for Amazon Location Maps for US East (N. Virginia) is: `https://maps.geo.us-east-1.amazonaws.com`.

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | **Maps:** maps.geo.us-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| | | **Places:** places.geo.us-east-2.amazonaws.com<br><br>**Routes:** routes.geo.us-east-2.amazonaws.com<br><br>**Geofences:** geofencing.geo.us-east-2.amazonaws.com<br><br>**Trackers:** tracking.geo.us-east-2.amazonaws.com<br><br>**Metadata:** metadata.geo.us-east-2.amazonaws.com | |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | **Maps:** maps.geo.us-east-1.amazonaws.com<br><br>**Places:** places.geo.us-east-1.amazonaws.com<br><br>**Routes:** routes.geo.us-east-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.us-east-1.amazonaws.com<br><br>**Trackers:** tracking.geo.us-east-1.amazonaws.com<br><br>**Metadata:** metadata.geo.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US West (Oregon) | us-west-2 | **Maps:** maps.geo.us-west-2.amazonaws.com<br><br>**Places:** places.geo.us-west-2.amazonaws.com<br><br>**Routes:** routes.geo.us-west-2.amazonaws.com<br><br>**Geofences:** geofencing.geo.us-west-2.amazonaws.com<br><br>**Trackers:** tracking.geo.us-west-2.amazonaws.com<br><br>**Metadata:** metadata.geo.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Malaysia) | ap-southeast-5 | **Maps:** maps.geo.ap-southeast-5.amazonaws.com<br><br>**Places:** places.geo.ap-southeast-5.amazonaws.com<br><br>**Routes:** routes.geo.ap-southeast-5.amazonaws.com<br><br>**Geofences:** geofencing.geo.ap-southeast-5.amazonaws.com<br><br>**Trackers:** tracking.geo.ap-southeast-5.amazonaws.com<br><br>**Metadata:** metadata.geo.ap-southeast-5.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Mumbai) | ap-south-1 | **Maps:** maps.geo.ap-south-1.amazonaws.com<br><br>**Places:** places.geo.ap-south-1.amazonaws.com<br><br>**Routes:** routes.geo.ap-south-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.ap-south-1.amazonaws.com<br><br>**Trackers:** tracking.geo.ap-south-1.amazonaws.com<br><br>**Metadata:** metadata.geo.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Singapore) | ap-southeast-1 | **Maps:** maps.geo.ap-southeast-1.amazonaws.com<br><br>**Places:** places.geo.ap-southeast-1.amazonaws.com<br><br>**Routes:** routes.geo.ap-southeast-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.ap-southeast-1.amazonaws.com<br><br>**Trackers:** tracking.geo.ap-southeast-1.amazonaws.com<br><br>**Metadata:** metadata.geo.ap-southeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Sydney) | ap-southeast-2 | **Maps:** maps.geo.ap-southeast-2.amazonaws.com<br><br>**Places:** places.geo.ap-southeast-2.amazonaws.com<br><br>**Routes:** routes.geo.ap-southeast-2.amazonaws.com<br><br>**Geofences:** geofencing.geo.ap-southeast-2.amazonaws.com<br><br>**Trackers:** tracking.geo.ap-southeast-2.amazonaws.com<br><br>**Metadata:** metadata.geo.ap-southeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Asia Pacific (Tokyo) | ap-northeast-1 | **Maps:** maps.geo.ap-northeast-1.amazonaws.com<br><br>**Places:** places.geo.ap-northeast-1.amazonaws.com<br><br>**Routes:** routes.geo.ap-northeast-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.ap-northeast-1.amazonaws.com<br><br>**Trackers:** tracking.geo.ap-northeast-1.amazonaws.com<br><br>**Metadata:** metadata.geo.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Canada (Central) | ca-central-1 | **Maps:** maps.geo.ca-central-1.amazonaws.com<br><br>**Places:** places.geo.ca-central-1.amazonaws.com<br><br>**Routes:** routes.geo.ca-central-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.ca-central-1.amazonaws.com<br><br>**Trackers:** tracking.geo.ca-central-1.amazonaws.com<br><br>**Metadata:** metadata.geo.ca-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Frankfurt) | eu-central-1 | **Maps:** maps.geo.eu-central-1.amazonaws.com<br><br>**Places:** places.geo.eu-central-1.amazonaws.com<br><br>**Routes:** routes.geo.eu-central-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.eu-central-1.amazonaws.com<br><br>**Trackers:** tracking.geo.eu-central-1.amazonaws.com<br><br>**Metadata:** metadata.geo.eu-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Ireland) | eu-west-1 | **Maps:** maps.geo.eu-west-1.amazonaws.com<br><br>**Places:** places.geo.eu-west-1.amazonaws.com<br><br>**Routes:** routes.geo.eu-west-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.eu-west-1.amazonaws.com<br><br>**Trackers:** tracking.geo.eu-west-1.amazonaws.com<br><br>**Metadata:** metadata.geo.eu-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (London) | eu-west-2 | **Maps:** maps.geo.eu-west-2.amazonaws.com<br><br>**Places:** places.geo.eu-west-2.amazonaws.com<br><br>**Routes:** routes.geo.eu-west-2.amazonaws.com<br><br>**Geofences:** geofencing.geo.eu-west-2.amazonaws.com<br><br>**Trackers:** tracking.geo.eu-west-2.amazonaws.com<br><br>**Metadata:** metadata.geo.eu-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Spain) | eu-south-2 | **Maps:** maps.geo.eu-south-2.amazonaws.com<br><br>**Places:** places.geo.eu-south-2.amazonaws.com<br><br>**Routes:** routes.geo.eu-south-2.amazonaws.com<br><br>**Geofences:** geofencing.geo.eu-south-2.amazonaws.com<br><br>**Trackers:** tracking.geo.eu-south-2.amazonaws.com<br><br>**Metadata:** metadata.geo.eu-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Stockholm) | eu-north-1 | **Maps:** maps.geo.eu-north-1.amazonaws.com<br><br>**Places:** places.geo.eu-north-1.amazonaws.com<br><br>**Routes:** routes.geo.eu-north-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.eu-north-1.amazonaws.com<br><br>**Trackers:** tracking.geo.eu-north-1.amazonaws.com<br><br>**Metadata:** metadata.geo.eu-north-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| South America (São Paulo) | sa-east-1 | **Maps:** maps.geo.sa-east-1.amazonaws.com<br><br>**Places:** places.geo.sa-east-1.amazonaws.com<br><br>**Routes:** routes.geo.sa-east-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.sa-east-1.amazonaws.com<br><br>**Trackers:** tracking.geo.sa-east-1.amazonaws.com<br><br>**Metadata:** metadata.geo.sa-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-West) | us-gov-west-1 | **Maps:** maps.geo.us-gov-west-1.amazonaws.com<br><br>**Places:** places.geo.us-gov-west-1.amazonaws.com<br><br>**Routes:** routes.geo.us-gov-west-1.amazonaws.com<br><br>**Geofences:** geofencing.geo.us-gov-west-1.amazonaws.com<br><br>**Trackers:** tracking.geo.us-gov-west-1.amazonaws.com<br><br>**Metadata:** metadata.geo.us-gov-west-1.amazonaws.com<br><br>**FIPS**<br><br>**Maps:** maps.geo-fips.us-gov-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| | | **Places:** places.geo-fips.us-gov-west-1.amazonaws.com<br><br>**Routes:** routes.geo-fips.us-gov-west-1.amazonaws.com<br><br>**Geofences:** geofencing.geo-fips.us-gov-west-1.amazonaws.com<br><br>**Trackers:** tracking.geo-fips.us-gov-west-1.amazonaws.com<br><br>**Metadata:** metadata.geo-fips.us-gov-west-1.amazonaws.com | |

## Control plane actions

The control plane actions for Amazon Location Service are:

| Service | Endpoint | API operation |
| --- | --- | --- |
| Amazon Location Maps | https://cp.maps.geo.*region*.api.aws | CreateMap |

| Service | Endpoint | API operation |
| --- | --- | --- |
| | https://cp.maps.geo.*region*.amazonaws.com | DeleteMap<br>DescribeMap<br>ListMaps<br>UpdateMap |
| Amazon Location Places | https://cp.places.geo.*region*.api.aws<br>https://cp.places.geo.*region*.amazonaws.com | CreatePlaceIndex<br>DeletePlaceIndex<br>DescribePlaceIndex<br>ListPlaceIndexes<br>UpdatePlaceIndex |
| Amazon Location Geofences | https://cp.geofencing.geo.*region*.api.aws<br>https://cp.geofencing.geo.*region*.amazonaws.com | CreateGeofenceCollection<br>DeleteGeofenceCollection<br>DescribeGeofenceCollection<br>ListGeofenceCollections<br>UpdateGeofenceCollection |

| Service | Endpoint | API operation |
| --- | --- | --- |
| Amazon Location Trackers | https://cp.tracking.geo.*region*.api.aws<br><br>https://cp.tracking.geo.*region*.amazonaws.com | CreateTracker<br><br>DeleteTracker<br><br>DescribeTracker<br><br>UpdateTracker<br><br>ListTrackers<br><br>AssociateTrackerConsumer<br><br>DisassociateTrackerConsumer<br><br>ListTrackerConsumers |
| Amazon Location Routes | https://cp.routes.geo.*region*.api.aws<br><br>https://cp.routes.geo.*region*.amazonaws.com | CreateRouteCalculator<br><br>DeleteRouteCalculator<br><br>DescribeRouteCalculator<br><br>ListRouteCalculators<br><br>UpdateRouteCalculator |
| Amazon Location Metadata | https://cp.metadata.geo.*region*.api.aws<br><br>https://cp.metadata.geo.*region*.amazonaws.com | CreateKey<br><br>DeleteKey<br><br>DescribeKey<br><br>ListKeys<br><br>UpdateKey |

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| API Key resources per account | Each supported Region: 500 | No | The maximum number of API key resources (active or expired) that you can have per account. |
| Geofence Collection resources per account | Each supported Region: 1,500 | Yes | The maximum number of Geofence Collection resources that you can create per account. |
| Geofences per Geofence Collection | Each supported Region: 50,000 | No | The maximum number of Geofences that you can create per Geofence Collection. |
| Map resources per account | Each supported Region: 40 | Yes | The maximum number of Map resources that you can create per account. |
| Place Index resources per account | Each supported Region: 40 | Yes | The maximum number of Place Index resources that you can create per account. |
| Rate of AssociateTrackerConsumer API requests | Each supported Region: 10 per second | Yes | The maximum number of AssociateTrackerConsumer requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of BatchDeleteDevicePositionHistory API requests | Each supported Region: 50 per second | Yes | The maximum number of BatchDeleteDevicePositionHistory requests that you can make per second. Additional requests are throttled. |
| Rate of BatchDeleteGeofence API requests | Each supported Region: 50 per second | Yes | The maximum number of BatchDeleteGeofence requests that you can make per second. Additional requests are throttled. |
| Rate of BatchEvaluateGeofences API requests | Each supported Region: 50 per second | Yes | The maximum number of BatchEvaluateGeofences requests that you can make per second. Additional requests are throttled. |
| Rate of BatchGetDevicePosition API requests | Each supported Region: 50 per second | Yes | The maximum number of BatchGetDevicePosition requests that you can make per second. Additional requests are throttled. |
| Rate of BatchPutGeofence API requests | Each supported Region: 50 per second | Yes | The maximum number of BatchPutGeofence requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of BatchUpdateDevicePosition API requests | Each supported Region: 50 per second | Yes | The maximum number of BatchUpdateDevicePosition requests that you can make per second. Additional requests are throttled. |
| Rate of CalculateRoute API requests | Each supported Region: 10 per second | Yes | The maximum number of CalculateRoute requests that you can make per second. Additional requests are throttled. |
| Rate of CalculateRouteMatrix API requests | Each supported Region: 5 per second | Yes | The maximum number of CalculateRouteMatrix requests that you can make per second. Additional requests are throttled. |
| Rate of CancelJob API requests | Each supported Region: 5 per second | Yes | The maximum number of CancelJob requests that you can make per second. Additional requests are throttled. |
| Rate of CreateGeofenceCollection API requests | Each supported Region: 10 per second | Yes | The maximum number of CreateGeofenceCollection requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateKey API requests | Each supported Region: 10 per second | Yes | The maximum number of CreateKey requests that you can make per second. Additional requests are throttled. |
| Rate of CreateMap API requests | Each supported Region: 10 per second | Yes | The maximum number of CreateMap requests that you can make per second. Additional requests are throttled. |
| Rate of CreatePlaceIndex API requests | Each supported Region: 10 per second | Yes | The maximum number of CreatePlaceIndex requests that you can make per second. Additional requests are throttled. |
| Rate of CreateRouteCalculator API requests | Each supported Region: 10 per second | Yes | The maximum number of CreateRouteCalculator requests that you can make per second. Additional requests are throttled. |
| Rate of CreateTracker API requests | Each supported Region: 10 per second | Yes | The maximum number of CreateTracker requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteGeofenceCollection API requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteGeofenceCollection requests that you can make per second. Additional requests are throttled. |
| Rate of DeleteKey API requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteKey requests that you can make per second. Additional requests are throttled. |
| Rate of DeleteMap API requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteMap requests that you can make per second. Additional requests are throttled. |
| Rate of DeletePlaceIndex API requests | Each supported Region: 10 per second | Yes | The maximum number of DeletePlaceIndex requests that you can make per second. Additional requests are throttled. |
| Rate of DeleteRouteCalculator API requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteRouteCalculator requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteTracker API requests | Each supported Region: 10 per second | Yes | The maximum number of DeleteTracker requests that you can make per second. Additional requests are throttled. |
| Rate of DescribeGeofenceCollection API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeGeofenceCollection requests that you can make per second. Additional requests are throttled. |
| Rate of DescribeKey API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeKey requests that you can make per second. Additional requests are throttled. |
| Rate of DescribeMap API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeMap requests that you can make per second. Additional requests are throttled. |
| Rate of DescribePlaceIndex API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribePlaceIndex requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeRouteCalculator API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeRouteCalculator requests that you can make per second. Additional requests are throttled. |
| Rate of DescribeTracker API requests | Each supported Region: 10 per second | Yes | The maximum number of DescribeTracker requests that you can make per second. Additional requests are throttled. |
| Rate of DisassociateTrackerConsumer API requests | Each supported Region: 10 per second | Yes | The maximum number of DisassociateTrackerConsumer requests that you can make per second. Additional requests are throttled. |
| Rate of ForecastGeofenceEvents API requests | Each supported Region: 50 per second | Yes | The maximum number of ForecastGeofenceEvents requests that you can make per second. Additional requests are throttled. |
| Rate of GetDevicePosition API requests | Each supported Region: 50 per second | Yes | The maximum number of GetDevicePosition requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetDevicePositionHistory API requests | Each supported Region: 50 per second | Yes | The maximum number of GetDevicePositionHistory requests that you can make per second. Additional requests are throttled. |
| Rate of GetGeofence API requests | Each supported Region: 50 per second | Yes | The maximum number of GetGeofence requests that you can make per second. Additional requests are throttled. |
| Rate of GetJob API requests | Each supported Region: 5 per second | Yes | The maximum number of GetJob requests that you can make per second. Additional requests are throttled. |
| Rate of GetMapGlyphs API requests | Each supported Region: 50 per second | Yes | The maximum number of GetMapGlyphs requests that you can make per second. Additional requests are throttled. |
| Rate of GetMapSprites API requests | Each supported Region: 50 per second | Yes | The maximum number of GetMapSprites requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetMapStyleDescriptor API requests | Each supported Region: 50 per second | Yes | The maximum number of GetMapStyleDescriptor requests that you can make per second. Additional requests are throttled. |
| Rate of GetMapTile API requests | Each supported Region: 500 per second | Yes | The maximum number of GetMapTile requests that you can make per second. Additional requests are throttled. |
| Rate of GetPlace API requests | Each supported Region: 100 per second | Yes | The maximum number of GetPlace requests that you can make per second. Additional requests are throttled. |
| Rate of ListDevicePositions API requests | Each supported Region: 50 per second | Yes | The maximum number of ListDevicePositions requests that you can make per second. Additional requests are throttled. |
| Rate of ListGeofenceCollections API requests | Each supported Region: 10 per second | Yes | The maximum number of ListGeofenceCollections requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListGeofences API requests | Each supported Region: 50 per second | Yes | The maximum number of ListGeofences requests that you can make per second. Additional requests are throttled. |
| Rate of ListJobs API requests | Each supported Region: 5 per second | Yes | The maximum number of ListJobs requests that you can make per second. Additional requests are throttled. |
| Rate of ListKeys API requests | Each supported Region: 10 per second | Yes | The maximum number of ListKeys requests that you can make per second. Additional requests are throttled. |
| Rate of ListMaps API requests | Each supported Region: 10 per second | Yes | The maximum number of ListMaps requests that you can make per second. Additional requests are throttled. |
| Rate of ListPlaceIndexes API requests | Each supported Region: 10 per second | Yes | The maximum number of ListPlaceIndexes requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListRouteCalculators API requests | Each supported Region: 10 per second | Yes | The maximum number of ListRouteCalculators requests that you can make per second. Additional requests are throttled. |
| Rate of ListTagsForResource API requests | Each supported Region: 10 per second | Yes | The maximum number of ListTagsForResource requests that you can make per second. Additional requests are throttled. |
| Rate of ListTrackerConsumers API requests | Each supported Region: 10 per second | Yes | The maximum number of ListTrackerConsumers requests that you can make per second. Additional requests are throttled. |
| Rate of ListTrackers API requests | Each supported Region: 10 per second | Yes | The maximum number of ListTrackers requests that you can make per second. Additional requests are throttled. |
| Rate of PutGeofence API requests | Each supported Region: 50 per second | Yes | The maximum number of PutGeofence requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of SearchPlaceIndexForPosition API requests | Each supported Region: 100 per second | Yes | The maximum number of SearchPlaceIndexForPosition requests that you can make per second. Additional requests are throttled. |
| Rate of SearchPlaceIndexForSuggestions API requests | Each supported Region: 100 per second | Yes | The maximum number of SearchPlaceIndexForSuggestions requests that you can make per second. Additional requests are throttled. |
| Rate of SearchPlaceIndexForText API requests | Each supported Region: 100 per second | Yes | The maximum number of SearchPlaceIndexForText requests that you can make per second. Additional requests are throttled. |
| Rate of StartJob API requests | Each supported Region: 5 per second | Yes | The maximum number of StartJob requests that you can make per second. Additional requests are throttled. |
| Rate of TagResource API requests | Each supported Region: 10 per second | Yes | The maximum number of TagResource requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UntagResource API requests | Each supported Region: 10 per second | Yes | The maximum number of UntagResource requests that you can make per second. Additional requests are throttled. |
| Rate of UpdateGeofenceCollection API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateGeofenceCollection requests that you can make per second. Additional requests are throttled. |
| Rate of UpdateKey API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateKey requests that you can make per second. Additional requests are throttled. |
| Rate of UpdateMap API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateMap requests that you can make per second. Additional requests are throttled. |
| Rate of UpdatePlaceIndex API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdatePlaceIndex requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of UpdateRouteCalculator API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateRouteCalculator requests that you can make per second. Additional requests are throttled. |
| Rate of UpdateTracker API requests | Each supported Region: 10 per second | Yes | The maximum number of UpdateTracker requests that you can make per second. Additional requests are throttled. |
| Rate of VerifyDevicePosition API requests | Each supported Region: 50 per second | Yes | The maximum number of VerifyDevicePosition requests that you can make per second. Additional requests are throttled. |
| Rate of geo-maps:GetStaticMap API requests | Each supported Region: 50 per second | Yes | The maximum number of geo-maps:GetStaticMap requests that you can make per second. Additional requests are throttled. |
| Rate of geo-maps:GetTile API requests | Each supported Region: 2,000 per second | Yes | The maximum number of geo-maps:GetTile requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of geo-places:Autocomplete API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:Autocomplete requests that you can make per second. Additional requests are throttled. |
| Rate of geo-places:Geocode API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:Geocode requests that you can make per second. Additional requests are throttled. |
| Rate of geo-places:GetPlace API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:GetPlace requests that you can make per second. Additional requests are throttled. |
| Rate of geo-places:ReverseGeocode API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:ReverseGeocode requests that you can make per second. Additional requests are throttled. |
| Rate of geo-places:SearchNearby API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:SearchNearby requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of geo-places:SearchText API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:SearchText requests that you can make per second. Additional requests are throttled. |
| Rate of geo-places:Suggest API requests | Each supported Region: 100 per second | Yes | The maximum number of geo-places:Suggest requests that you can make per second. Additional requests are throttled. |
| Rate of geo-routes:CalculateIsolines API requests | Each supported Region: 20 per second | Yes | The maximum number of geo-routes:CalculateIsolines requests that you can make per second. Additional requests are throttled. |
| Rate of geo-routes:CalculateRouteMatrix API requests | Each supported Region: 5 per second | Yes | The maximum number of geo-routes:CalculateRouteMatrix requests that you can make per second. Additional requests are throttled. |
| Rate of geo-routes:CalculateRoutes API requests | Each supported Region: 20 per second | Yes | The maximum number of geo-routes:CalculateRoutes requests that you can make per second. Additional requests are throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of geo-routes:OptimizeWaypoints API requests | Each supported Region: 5 per second | Yes | The maximum number of geo-routes:OptimizeWaypoints requests that you can make per second. Additional requests are throttled. |
| Rate of geo-routes:SnapToRoads API requests | Each supported Region: 20 per second | Yes | The maximum number of geo-routes:SnapToRoads requests that you can make per second. Additional requests are throttled. |
| Route Calculator resources per account | Each supported Region: 40 | Yes | The maximum number of Route Calculator resources that you can create per account. |
| Tracker consumers per tracker | Each supported Region: 5 | No | The maximum number of Geofence Collection that Tracker resource can be associated with. |
| Tracker resources per account | Each supported Region: 500 | Yes | The maximum number of Tracker resources that you can create per account. |

For more information, see Amazon Location Service Quotas in the *Amazon Location Service Developer Guide*.

# Amazon Lookout for Equipment endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | lookoutequipment.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | lookoutequipment.ap-northeast-2.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | lookoutequipment.eu-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Components per dataset | Each supported Region: 3,000 | No | Maximum number of components per dataset. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Datasets | Each supported Region: 15 | Yes | Maximum number of datasets per account. |
| Inference schedulers per model | Each supported Region: 1 | No | Maximum number of inference schedulers per model. |
| Label groups | Each supported Region: 15 | Yes | Maximum number of label groups per account. |
| Length of component name | Each supported Region: 200 | No | Maximum length of component name. |
| Models | Each supported Region: 15 | Yes | Maximum number of models per account. |
| Number of columns across components in training data (excluding timestamp) | Each supported Region: 300 | No | Maximum number of columns across components in training data (excluding timestamp) |
| Number of columns across components per dataset (excluding timestamp) | Each supported Region: 3,000 | No | Maximum number of columns across components per dataset (excluding timestamp). |
| Number of components in training data | Each supported Region: 300 | No | Maximum number of components in training data. |
| Number of files per component (per dataset) | Each supported Region: 1,000 | No | Maximum number of files per component (per dataset). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of files per component (per inference execution) | Each supported Region: 60 | No | Maximum number of files per component (per inference execution). |
| Number of import dataset jobs in progress | Each supported Region: 15 | Yes | Maximum number of import dataset jobs in progress per account. |
| Number of import model jobs in progress | Each supported Region: 15 | Yes | Maximum number of import model jobs in progress per account. |
| Number of labels per label group | Each supported Region: 5,000 | Yes | Maximum number of labels per label group. |
| Number of rows in evaluation data (after resampling) | Each supported Region: 1,500,000 | No | Maximum number of rows in evaluation data (after resampling). |
| Number of rows in inference input data, after resampling (1-hour scheduling frequency) | Each supported Region: 3,600 | No | Maximum number of rows in inference input data, after resampling (1-hour scheduling frequency). |
| Number of rows in inference input data, after resampling (10-min scheduling frequency) | Each supported Region: 600 | No | Maximum number of rows in inference input data, after resampling (10-min scheduling frequency). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of rows in inference input data, after resampling (15-min scheduling frequency) | Each supported Region: 900 | No | Maximum number of rows in inference input data, after resampling (15-min scheduling frequency). |
| Number of rows in inference input data, after resampling (30-min scheduling frequency) | Each supported Region: 1,800 | No | Maximum number of rows in inference input data, after resampling (30-min scheduling frequency). |
| Number of rows in inference input data, after resampling (5-min scheduling frequency) | Each supported Region: 300 | No | Maximum number of rows in inference input data, after resampling (5-min scheduling frequency). |
| Number of rows in training data (after resampling) | Each supported Region: 1,500,000 | No | Maximum number of rows in training data (after resampling). |
| Pending data ingestion jobs | Each supported Region: 5 | Yes | Maximum number of pending data ingestion jobs per account. |
| Pending models | Each supported Region: 5 | Yes | Maximum number of pending models per account. |
| Size of raw data in inference input data (1-hour scheduling frequency) | Each supported Region: 60 Megabytes | No | Maximum size of raw data in inference input data (1-hour scheduling frequency). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Size of raw data in inference input data (10-min scheduling frequency) | Each supported Region: 10 Megabytes | No | Maximum size of raw data in inference input data (10-min scheduling frequency). |
| Size of raw data in inference input data (15-min scheduling frequency) | Each supported Region: 15 Megabytes | No | Maximum size of raw data in inference input data (15-min scheduling frequency). |
| Size of raw data in inference input data (30-min scheduling frequency) | Each supported Region: 30 Megabytes | No | Maximum size of raw data in inference input data (30-min scheduling frequency). |
| Size of raw data in inference input data (5-min scheduling frequency) | Each supported Region: 5 Megabytes | No | Maximum size of raw data in inference input data (5-min scheduling frequency). |
| Size per dataset | Each supported Region: 50 Gigabytes | No | Maximum size per dataset. |
| Size per file | Each supported Region: 5 Gigabytes | No | Maximum size per file. |
| Timespan of training data | Each supported Region: 180 per day | No | Minimum timespan of training data |

# Amazon Macie endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | macie2.us-east-2.amazonaws.com | HTTPS | |
| | | macie2-fips.us-east-2.api.aws | HTTPS | |
| | | macie2-fips.us-east-2.amazonaws.com | HTTPS | |
| | | macie2.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | macie2.us-east-1.amazonaws.com | HTTPS | |
| | | macie2-fips.us-east-1.api.aws | HTTPS | |
| | | macie2-fips.us-east-1.amazonaws.com | HTTPS | |
| | | macie2.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | macie2.us-west-1.amazonaws.com | HTTPS | |
| | | macie2-fips.us-west-1.api.aws | HTTPS | |
| | | macie2-fips.us-west-1.amazonaws.com | HTTPS | |
| | | macie2.us-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | macie2.us-west-2.amazonaws.com | HTTPS |
| | | macie2-fips.us-west-2.api.aws | HTTPS |
| | | macie2-fips.us-west-2.amazonaws.com | HTTPS |
| | | macie2.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | macie2.af-south-1.amazonaws.com | HTTPS |
| | | macie2.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | macie2.ap-east-1.amazonaws.com | HTTPS |
| | | macie2.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | macie2.ap-south-1.amazonaws.com | HTTPS |
| | | macie2.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | macie2.ap-northeast-3.amazonaws.com | HTTPS |
| | | macie2.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | macie2.ap-northeast-2.amazonaws.com | HTTPS |
| | | macie2.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | macie2.ap-southeast-1.amazonaws.com | HTTPS |
| | | macie2.ap-southeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | macie2.ap-southeast-2.amazonaws.com | HTTPS | |
| | | macie2.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | macie2.ap-northeast-1.amazonaws.com | HTTPS | |
| | | macie2.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | macie2.ca-central-1.amazonaws.com | HTTPS | |
| | | macie2.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | macie2.eu-central-1.amazonaws.com | HTTPS | |
| | | macie2.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | macie2.eu-west-1.amazonaws.com | HTTPS | |
| | | macie2.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | macie2.eu-west-2.amazonaws.com | HTTPS | |
| | | macie2.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | macie2.eu-south-1.amazonaws.com | HTTPS | |
| | | macie2.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | macie2.eu-west-3.amazonaws.com | HTTPS | |
| | | macie2.eu-west-3.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | macie2.eu-north-1.amazonaws.com | HTTPS | |
| | | macie2.eu-north-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | macie2.il-central-1.amazonaws.com<br>macie2.il-central-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | macie2.me-south-1.amazonaws.com<br>macie2.me-south-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | macie2.sa-east-1.amazonaws.com<br>macie2.sa-east-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Apache Avro container (.avro) file size | Each supported Region: 8 Gigabytes | No | The maximum size (in GB) of an individual Apache Avro object container (.avro) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |
| Apache Parquet (.parquet) file size | Each supported Region: 8 Gigabytes | No | The maximum size (in GB) of an individual Apache Parquet (.parquet) file that Macie can analyze. If a file is larger, |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | Macie doesnt analyze any data in the file. |
| Custom data identifiers per account | Each supported Region: 10,000 | No | The maximum number of custom data identifiers that can be created for this account in the current Region. |
| Custom data identifiers per sensitive data discovery job | Each supported Region: 30 | No | The maximum number of custom data identifiers that you can configure a sensitive data discovery job to use. |
| Extracted archive bytes | Each supported Region: 10 Gigabytes | No | The maximum amount of data (in GB) that Macie can extract and analyze in a compressed or archive file. If Macie starts analyzing this type of file and determines that the file contains more than this amount of data, Macie stops analyzing the file and creates sensitive data findings and discovery results only for the data that was processed. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Extracted archive files | Each supported Region: 1,000,000 | No | The maximum number of files that Macie can extract and analyze in an archive file. If Macie starts analyzing data in an archive file and determines that the file contains more than the maximum number of files, Macie stops analyzing data in the file and creates sensitive data findings and discovery results only for the data that was processed. |
| Findings rules | Each supported Region: 1,000 | No | The maximum number of filter rules and suppression rules that you can create for this account in the current Region. |
| Full names detected | Each supported Region: 1,000 | No | The maximum number of full names that Macie can detect and report for a file, including individual archive files. After Macie detects the maximum number, Macie stops incrementing the count and reporting location data for full names. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GNU Zip compressed archive (.gz or .gzip) file size | Each supported Region: 8 Gigabytes | No | The maximum size (in GB) of an individual GNU Zip compressed archive (.gz or .gzip) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |
| Mailing addresses detected | Each supported Region: 1,000 | No | The maximum number of mailing addresses that Macie can detect and report for a file, including individual archive files. After Macie detects the maximum number, Macie stops incrementing the count and reporting location data for mailing addresses. |
| Member accounts by invitation | Each supported Region: 1,000 | No | The maximum number of member accounts that can be associated with a Macie administrator account by invitation in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Member accounts through AWS Organizations | Each supported Region: 5,000 | No | The maximum number of member accounts that can be associated with the Macie administrator account for an AWS Organizations organization in the current Region. |
| Microsoft Excel workbook (.xls or .xlsx) file size | Each supported Region: 512 Megabytes | No | The maximum size (in MB) of an individual Microsoft Excel workbook (.xls or .xlsx) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |
| Microsoft Word document (.doc or .docx) file size | Each supported Region: 512 Megabytes | No | The maximum size (in MB) of an individual Microsoft Word document (.doc or .docx) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Nested levels | Each supported Region: 10 | No | The maximum number of nested levels that Macie can analyze in an archive file. If the metadata for an archive file indicates that the file contains more than the maximum number of nested levels, Macie doesnt extract or analyze any data in the file. |
| Nested levels in structured data | Each supported Region: 256 | No | The maximum number of nested levels that Macie can analyze in a JSON (.json) or JSON Lines (.jsonl) file. If a file contains more than the maximum number of nested levels, Macie doesnt analyze any data in the file. |
| Non-binary text file size | Each supported Region: 20 Gigabytes | No | The maximum size (in GB) of an individual non-binary text file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Portable Document Format (.pdf) file size | Each supported Region: 1,024 Megabytes | No | The maximum size (in MB) of an individual Portable Document Format (.pdf) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |
| Preventative control monitoring | Each supported Region: 10,000 | No | The maximum number of S3 buckets that Macie monitors for security and access control. |
| S3 buckets per sensitive data discovery job | Each supported Region: 1,000 | No | The maximum number of S3 buckets that you can explicitly select for a sensitive data discovery job to analyze. If youre the Macie administrator for an organization, the buckets can span as many as 1,000 accounts in your organization. |
| Sensitive data discovery occurrences | Each supported Region: 1,000 | No | The maximum number of occurrences of each type of sensitive data that Macie detects and provides detailed location data for in sensitive data discovery results. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Sensitive data discovery per month per account | Each supported Region: 5 Terabytes | Yes | The maximum amount of data (in TB) that you can analyze by running sensitive data discovery jobs for this account in the current Region. |
| Sensitive data finding occurrences | Each supported Region: 15 | No | The maximum number of detection locations that Macie provides in a sensitive data finding. |
| TAR archive (.tar) file size | Each supported Region: 20 Gigabytes | No | The maximum size (in GB) of an individual TAR archive (.tar) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |
| ZIP compressed archive (.zip) file size | Each supported Region: 8 Gigabytes | No | The maximum size (in GB) of an individual ZIP compressed archive (.zip) file that Macie can analyze. If a file is larger, Macie doesnt analyze any data in the file. |

For more information, see Amazon Macie quotas in the *Amazon Macie User Guide*.

## AWS Mainframe Modernization endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | m2.us-east-2.amazonaws.com | HTTPS |
| | | m2-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | m2.us-east-1.amazonaws.com | HTTPS |
| | | m2-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | m2.us-west-1.amazonaws.com | HTTPS |
| | | m2-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | m2.us-west-2.amazonaws.com | HTTPS |
| | | m2-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | m2.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | m2.ap-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | m2.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | m2.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | m2.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | m2.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | m2.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | m2.ca-central-1.amazonaws.com<br>m2-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | m2.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | m2.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | m2.eu-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Milan) | eu-south-1 | m2.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | m2.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | m2.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | m2.eu-north-1.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | m2.il-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | m2.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | m2.us-gov-east-1.amazonaws.com<br>m2-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | m2.us-gov-west-1.amazonaws.com<br>m2-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Max Applications Per AWS Account | Each supported Region: 20 | Yes | Max Applications Per AWS Account |
| Max DataTransferEndpoints Per AWS Account | Each supported Region: 10 | No | Max DataTransferEndpoints Per AWS Account |
| Max Environments Per AWS Account | Each supported Region: 20 | Yes | Max Environments Per AWS Account |
| Max Instances Per High Availability Environment | Each supported Region: 12 | Yes | Max Instances Per High Availability Environment |
| Max number of EFS filesystems per environment | Each supported Region: 1 | Yes | Max number of EFS filesystems per environment |
| Max number of FSX filesystems per environment | Each supported Region: 1 | Yes | Max number of FSX filesystems per environment |

# Amazon Machine Learning endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | machinelearning.us-east-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | machinelearning.eu-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Batch prediction input records | Each supported Region: 100,000,000 | Yes | The maximum number of records of batch prediction input. |
| Batch prediction input size | Each supported Region: 1 Terabytes | Yes | The maximum size (in TB) of batch prediction input. |
| Classes for multiclass ML models | Each supported Region: 100 | Yes | The maximum number of classes for multiclass ML models. |
| Job runtime | Each supported Region: 7 | No | The maximum runtime length (in days) for any job. |
| ML model size | Each supported Region: 2 Gigabytes | No | The maximum ML model size (in GB). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Observation size | Each supported Region: 100 Kilobytes | Yes | The maximum size (in KB) of each observation. |
| Rate of real-time prediction requests per endpoint | Each supported Region: 200 | Yes | The maximum number of requests per second that you can perform with each real-time prediction endpoint. |
| Recipe complexity | Each supported Region: 10,000 | Yes | The maximum recipe complexity (number of processed output variables). |
| Simultaneous jobs | Each supported Region: 25 | Yes | The maximum number of simultaneous jobs. |
| Tags per object | Each supported Region: 50 | No | The maximum number of tags per object. |
| Total RAM for all real-time prediction endpoints | Each supported Region: 10 Gigabytes | Yes | The maximum total RAM (in GB) for all real-time prediction endpoints. |
| Total rate of all real-time prediction requests | Each supported Region: 10,000 | Yes | The maximum total number of requests per second that you can perform with all of your real-time prediction endpoints. |
| Training data size | Each supported Region: 100 Gigabytes | Yes | The maximum size (in GB) of training data. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Variables per data file | Each supported Region: 1,000 | Yes | The maximum number of variables in a data file (schema). |

For more information, see Amazon ML Quotas in the *Amazon Machine Learning Developer Guide*.

## Amazon Managed Blockchain endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

For information about attributes of Starter Edition and Standard Edition networks, such as the number of members per network, peer nodes per member, available instance types, and more, see Amazon Managed Blockchain Pricing.

**Topics**

- Amazon Managed Blockchain (AMB) Access service endpoints and quotas
- Amazon Managed Blockchain (AMB) Query service endpoints and quotas

# Amazon Managed Blockchain (AMB) Access service endpoints and quotas

## AMB Access Regions and endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | managedblockchain.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | managedblockchain.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | managedblockchain.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | managedblockchain.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | managedblockchain.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | managedblockchain.eu-west-2.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | managedblockchain.us-gov-west-1.amazonaws.com | HTTPS | |

## AMB Access quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Number of Ethereum nodes an AWS account can have | Each supported Region: 50 | Yes | The maximum number of Ethereum nodes an AWS account can have. |
| Number of Hyperledger Fabric channels per Standard Edition network | Each supported Region: 8 | No | The maximum number of Hyperledger Fabric channels per Standard Edition network. |
| Number of Hyperledger Fabric channels per Starter Edition network | Each supported Region: 8 | No | The maximum number of Hyperledger Fabric channels per Starter Edition network. |
| Number of Standard Edition networks in which an AWS account can have a member | Each supported Region: 6 | Yes | The maximum number of Hyperledger Fabric Standard Edition networks in which an AWS account can have a member. |
| Number of Starter Edition networks in which an AWS account can have a member | Each supported Region: 6 | Yes | The maximum number of Hyperledger Fabric Starter Edition networks in which an AWS account can have a member. |

# Amazon Managed Blockchain (AMB) Query service endpoints and quotas

## AMB Query Regions and endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | managedblockchain-query.us-east-1.amazonaws.com | HTTPS | |

## AMB Query quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of BatchGetTokenBalance requests per second | Each supported Region: 10 per second | Yes | The maximum number of BatchGetTokenBalance requests you can make per second per account. |
| Number of GetAssetContract requests per second | Each supported Region: 100 per second | Yes | The maximum number of GetAssetContract requests you can make per second per account. |
| Number of GetTokenBalance requests per second | Each supported Region: 100 per second | Yes | The maximum number of GetTokenBalance requests you can make per second per account. |
| Number of GetTransaction requests per second | Each supported Region: 100 per second | Yes | The maximum number of GetTransaction requests you can make per second per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ListAssetContracts requests per second | Each supported Region: 100 per second | Yes | The maximum number of ListAssetContracts requests you can make per second per account. |
| Number of ListFilteredTransactionEvents requests per second | Each supported Region: 100 per second | Yes | The maximum number of ListFilteredTransactionEvents requests you can make per second per account. |
| Number of ListTokenBalances requests per second | Each supported Region: 100 per second | Yes | The maximum number of ListTokenBalances requests you can make per second per account. |
| Number of ListTransactionEvents requests per second | Each supported Region: 200 per second | Yes | The maximum number of ListTransactionEvents requests you can make per second per account. |
| Number of ListTransactions requests per second | Each supported Region: 200 per second | Yes | The maximum number of ListTransactions requests you can make per second per account. |

# Amazon Managed Service for Apache Flink endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | kinesisanalytics.us-east-2.amazonaws.com<br>kinesisanalytics-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | kinesisanalytics.us-east-1.amazonaws.com<br>kinesisanalytics-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | kinesisanalytics.us-west-1.amazonaws.com<br>kinesisanalytics-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | kinesisanalytics.us-west-2.amazonaws.com<br>kinesisanalytics-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | kinesisanalytics.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | kinesisanalytics.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | kinesisanalytics.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | kinesisanalytics.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | kinesisanalytics.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | kinesisanalytics.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | kinesisanalytics.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | kinesisanalytics.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | kinesisanalytics.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | kinesisanalytics.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | kinesisanalytics.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | kinesisanalytics.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | kinesisanalytics.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | kinesisanalytics.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | kinesisanalytics.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | kinesisanalytics.ca-central-1.amazonaws.com<br>kinesisanalytics-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | kinesisanalytics.ca-west-1.amazonaws.com<br>kinesisanalytics-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | kinesisanalytics.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | kinesisanalytics.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | kinesisanalytics.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | kinesisanalytics.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | kinesisanalytics.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | kinesisanalytics.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | kinesisanalytics.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | kinesisanalytics.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | kinesisanalytics.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | kinesisanalytics.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | kinesisanalytics.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | kinesisanalytics.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | kinesisanalytics.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | kinesisanalytics.us-gov-east-1.amazonaws.com<br><br>kinesisanalytics-fips.us-gov-east-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | kinesisanalytics.us-gov-west-1.amazonaws.com<br><br>kinesisanalytics-fips.us-gov-west-1.amazonaws.com | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Apache Flink Kinesis Processing Units (KPUs) | Each supported Region: 64 | Yes | The maximum number of Kinesis Processing Units (KPUs) that your Apache Flink application can use. |
| Application count | Each supported Region: 100 | Yes | The maximum number of applications per account per Region. |
| Input Parallelism in input streams for SQL applications | Each supported Region: 64 | No | The maximum number of in-application input |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| | | | streams for SQL applications. |
| SQL Kinesis Processing Units (KPUs) | Each supported Region: 8 | Yes | The maximum number of Kinesis Processing Units (KPUs) that your SQL application can use. |

For more information, see Quotas in the *Amazon Managed Service for Apache Flink for Apache Flink Developer Guide*.

# AWS Managed Services endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

See AMS VPC endpoints and AMS Supported configurations in the *AWS Managed Services User Guide*.

## Service quotas

See AMS account limits in the *AWS Managed Services User Guide*.

# AWS Management Console endpoints and quotas

AWS Management Console has Regional endpoints that allow you to directly access the console in a given AWS Region. The general syntax of a Regional endpoint is as follows:

```
https://region-code.console.aws.amazon.com
```

For example, `https://us-west-2.console.aws.amazon.com` is the endpoint for the AWS Management Console service in the US West (Oregon) Region.

The table below lists the name, code, and endpoint of each AWS Region.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | us-east-2.console.aws.amazon.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | us-east-1.console.aws.amazon.com | HTTPS | |
| US West (N. California) | us-west-1 | us-west-1.console.aws.amazon.com | HTTPS | |
| US West (Oregon) | us-west-2 | us-west-2.console.aws.amazon.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | af-south-1.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | ap-east-1.console.aws.amazon.com | HTTPS | |
| Asia Pacific | ap-south-2 | ap-south-2.console.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ap-southeast-3.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | ap-southeast-5.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ap-southeast-4.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | ap-south-1.console.aws.amazon.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ap-southeast-6.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ap-northeast-3.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ap-northeast-2.console.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | ap-southeast-1.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ap-southeast-2.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | ap-east-2.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | ap-southeast-7.console.aws.amazon.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | ap-northeast-1.console.aws.amazon.com | HTTPS | |
| Canada (Central) | ca-central-1 | ca-central-1.console.aws.amazon.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | ca-west-1.console.aws.amazon.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | eu-central-1.console.aws.amazon.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | eu-west-1.console.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | eu-west-2.console.aws.amazon.com | HTTPS | |
| Europe (Milan) | eu-south-1 | eu-south-1.console.aws.amazon.com | HTTPS | |
| Europe (Paris) | eu-west-3 | eu-west-3.console.aws.amazon.com | HTTPS | |
| Europe (Spain) | eu-south-2 | eu-south-2.console.aws.amazon.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | eu-north-1.console.aws.amazon.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | eu-central-2.console.aws.amazon.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | il-central-1.console.aws.amazon.com | HTTPS | |
| Mexico (Central) | mx-central-1 | mx-central-1.console.aws.amazon.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | me-south-1.console.aws.amazon.com | HTTPS | |
| Middle East (UAE) | me-central-1 | me-central-1.console.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | sa-east-1.console.aws.amazon.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | us-gov-east-1.console.amazonaws-us-gov.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | us-gov-west-1.console.amazonaws-us-gov.com | HTTPS | |

### Service quotas

This service has no quotas.

## Amazon Managed Workflows for Apache Airflow endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | api.airflow.us-east-2.amazonaws.com | HTTPS | |
| | | env.airflow.us-east-2.api.aws | HTTPS | |
| | | ops.airflow.us-east-2.amazonaws.com | HTTPS | |
| | | api.airflow.us-east-2.api.aws | HTTPS | |
| | | env.airflow-fips.us-east-2.api.aws | HTTPS | |
| | | env.airflow-fips.us-east-2.amazonaws.com | HTTPS | |
| | | api.airflow-fips.us-east-2.api.aws | HTTPS | |
| | | api.airflow-fips.us-east-2.amazonaws.com | HTTPS | |
| | | env.airflow.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | api.airflow.us-east-1.amazonaws.com | HTTPS | |
| | | env.airflow.us-east-1.api.aws | HTTPS | |
| | | ops.airflow.us-east-1.amazonaws.com | HTTPS | |
| | | api.airflow.us-east-1.api.aws | HTTPS | |
| | | env.airflow-fips.us-east-1.api.aws | HTTPS | |
| | | env.airflow-fips.us-east-1.amazonaws.com | HTTPS | |
| | | api.airflow-fips.us-east-1.api.aws | HTTPS | |
| | | api.airflow-fips.us-east-1.amazonaws.com | HTTPS | |
| | | env.airflow.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (N. California) | us-west-1 | api.airflow.us-west-1.amazonaws.com | HTTPS | |
| | | env.airflow.us-west-1.api.aws | HTTPS | |
| | | api.airflow.us-west-1.api.aws | HTTPS | |
| | | env.airflow-fips.us-west-1.api.aws | HTTPS | |
| | | env.airflow-fips.us-west-1.amazonaws.com | HTTPS | |
| | | api.airflow-fips.us-west-1.api.aws | HTTPS | |
| | | api.airflow-fips.us-west-1.amazonaws.com | HTTPS | |
| | | env.airflow.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | api.airflow.us-west-2.amazonaws.com | HTTPS | |
| | | env.airflow.us-west-2.api.aws | HTTPS | |
| | | ops.airflow.us-west-2.amazonaws.com | HTTPS | |
| | | api.airflow.us-west-2.api.aws | HTTPS | |
| | | env.airflow-fips.us-west-2.api.aws | HTTPS | |
| | | env.airflow-fips.us-west-2.amazonaws.com | HTTPS | |
| | | api.airflow-fips.us-west-2.api.aws | HTTPS | |
| | | api.airflow-fips.us-west-2.amazonaws.com | HTTPS | |
| | | env.airflow.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | api.airflow.af-south-1.amazonaws.com | HTTPS | |
| | | env.airflow.af-south-1.api.aws | HTTPS | |
| | | api.airflow.af-south-1.api.aws | HTTPS | |
| | | env.airflow.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | api.airflow.ap-east-1.amazonaws.com | HTTPS | |
| | | env.airflow.ap-east-1.api.aws | HTTPS | |
| | | api.airflow.ap-east-1.api.aws | HTTPS | |
| | | env.airflow.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | api.airflow.ap-south-2.amazonaws.com | HTTPS | |
| | | env.airflow.ap-south-2.api.aws | HTTPS | |
| | | api.airflow.ap-south-2.api.aws | HTTPS | |
| | | env.airflow.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | api.airflow.ap-southeast-3.amazonaws.com | HTTPS | |
| | | env.airflow.ap-southeast-3.api.aws | HTTPS | |
| | | api.airflow.ap-southeast-3.api.aws | HTTPS | |
| | | env.airflow.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | api.airflow.ap-southeast-5.amazonaws.com | HTTPS | |
| | | env.airflow.ap-southeast-5.api.aws | HTTPS | |
| | | api.airflow.ap-southeast-5.api.aws | HTTPS | |
| | | env.airflow.ap-southeast-5.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | api.airflow.ap-southeast-4.amazonaws.com | HTTPS |
| | | env.airflow.ap-southeast-4.api.aws | HTTPS |
| | | api.airflow.ap-southeast-4.api.aws | HTTPS |
| | | env.airflow.ap-southeast-4.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | api.airflow.ap-south-1.amazonaws.com | HTTPS |
| | | env.airflow.ap-south-1.api.aws | HTTPS |
| | | ops.airflow.ap-south-1.amazonaws.com | HTTPS |
| | | api.airflow.ap-south-1.api.aws | HTTPS |
| | | env.airflow.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | api.airflow.ap-northeast-3.amazonaws.com | HTTPS |
| | | env.airflow.ap-northeast-3.api.aws | HTTPS |
| | | api.airflow.ap-northeast-3.api.aws | HTTPS |
| | | env.airflow.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | api.airflow.ap-northeast-2.amazonaws.com | HTTPS |
| | | env.airflow.ap-northeast-2.api.aws | HTTPS |
| | | ops.airflow.ap-northeast-2.amazonaws.com | HTTPS |
| | | api.airflow.ap-northeast-2.api.aws | HTTPS |
| | | env.airflow.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | api.airflow.ap-southeast-1.amazonaws.com | HTTPS |
| | | env.airflow.ap-southeast-1.api.aws | HTTPS |
| | | ops.airflow.ap-southeast-1.amazonaws.com | HTTPS |
| | | api.airflow.ap-southeast-1.api.aws | HTTPS |
| | | env.airflow.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | api.airflow.ap-southeast-2.amazonaws.com | HTTPS |
| | | env.airflow.ap-southeast-2.api.aws | HTTPS |
| | | ops.airflow.ap-southeast-2.amazonaws.com | HTTPS |
| | | api.airflow.ap-southeast-2.api.aws | HTTPS |
| | | env.airflow.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | api.airflow.ap-southeast-7.amazonaws.com | HTTPS |
| | | env.airflow.ap-southeast-7.api.aws | HTTPS |
| | | api.airflow.ap-southeast-7.api.aws | HTTPS |
| | | env.airflow.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.airflow.ap-northeast-1.amazonaws.com | HTTPS |
| | | env.airflow.ap-northeast-1.api.aws | HTTPS |
| | | ops.airflow.ap-northeast-1.amazonaws.com | HTTPS |
| | | api.airflow.ap-northeast-1.api.aws | HTTPS |
| | | env.airflow.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | api.airflow.ca-central-1.amazonaws.com | HTTPS | |
| | | env.airflow.ca-central-1.api.aws | HTTPS | |
| | | ops.airflow.ca-central-1.amazonaws.com | HTTPS | |
| | | api.airflow.ca-central-1.api.aws | HTTPS | |
| | | env.airflow-fips.ca-central-1.api.aws | HTTPS | |
| | | env.airflow-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | api.airflow-fips.ca-central-1.api.aws | HTTPS | |
| | | api.airflow-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | env.airflow.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | api.airflow.ca-west-1.amazonaws.com | HTTPS | |
| | | env.airflow.ca-west-1.api.aws | HTTPS | |
| | | api.airflow.ca-west-1.api.aws | HTTPS | |
| | | env.airflow-fips.ca-west-1.api.aws | HTTPS | |
| | | env.airflow-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | api.airflow-fips.ca-west-1.api.aws | HTTPS | |
| | | api.airflow-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | env.airflow.ca-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | api.airflow.eu-central-1.amazonaws.com | HTTPS |
| | | env.airflow.eu-central-1.api.aws | HTTPS |
| | | ops.airflow.eu-central-1.amazonaws.com | HTTPS |
| | | api.airflow.eu-central-1.api.aws | HTTPS |
| | | env.airflow.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | api.airflow.eu-west-1.amazonaws.com | HTTPS |
| | | env.airflow.eu-west-1.api.aws | HTTPS |
| | | ops.airflow.eu-west-1.amazonaws.com | HTTPS |
| | | api.airflow.eu-west-1.api.aws | HTTPS |
| | | env.airflow.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | api.airflow.eu-west-2.amazonaws.com | HTTPS |
| | | env.airflow.eu-west-2.api.aws | HTTPS |
| | | ops.airflow.eu-west-2.amazonaws.com | HTTPS |
| | | api.airflow.eu-west-2.api.aws | HTTPS |
| | | env.airflow.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | api.airflow.eu-south-1.amazonaws.com | HTTPS |
| | | env.airflow.eu-south-1.api.aws | HTTPS |
| | | api.airflow.eu-south-1.api.aws | HTTPS |
| | | env.airflow.eu-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | api.airflow.eu-west-3.amazonaws.com | HTTPS |
| | | env.airflow.eu-west-3.api.aws | HTTPS |
| | | ops.airflow.eu-west-3.amazonaws.com | HTTPS |
| | | api.airflow.eu-west-3.api.aws | HTTPS |
| | | env.airflow.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | api.airflow.eu-south-2.amazonaws.com | HTTPS |
| | | env.airflow.eu-south-2.api.aws | HTTPS |
| | | api.airflow.eu-south-2.api.aws | HTTPS |
| | | env.airflow.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | api.airflow.eu-north-1.amazonaws.com | HTTPS |
| | | env.airflow.eu-north-1.api.aws | HTTPS |
| | | ops.airflow.eu-north-1.amazonaws.com | HTTPS |
| | | api.airflow.eu-north-1.api.aws | HTTPS |
| | | env.airflow.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | api.airflow.eu-central-2.amazonaws.com | HTTPS |
| | | env.airflow.eu-central-2.api.aws | HTTPS |
| | | api.airflow.eu-central-2.api.aws | HTTPS |
| | | env.airflow.eu-central-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | api.airflow.il-central-1.amazonaws.com | HTTPS |
| | | env.airflow.il-central-1.api.aws | HTTPS |
| | | api.airflow.il-central-1.api.aws | HTTPS |
| | | env.airflow.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | api.airflow.me-south-1.amazonaws.com | HTTPS |
| | | env.airflow.me-south-1.api.aws | HTTPS |
| | | api.airflow.me-south-1.api.aws | HTTPS |
| | | env.airflow.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | api.airflow.me-central-1.amazonaws.com | HTTPS |
| | | env.airflow.me-central-1.api.aws | HTTPS |
| | | api.airflow.me-central-1.api.aws | HTTPS |
| | | env.airflow.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | api.airflow.sa-east-1.amazonaws.com | HTTPS |
| | | env.airflow.sa-east-1.api.aws | HTTPS |
| | | ops.airflow.sa-east-1.amazonaws.com | HTTPS |
| | | api.airflow.sa-east-1.api.aws | HTTPS |
| | | env.airflow.sa-east-1.amazonaws.com | HTTPS |

**Note**

For IPv6, use `*.api.aws` in your endpoint, not `*.amazonaws.com`, for all regions except `ap-southeast-5`. For example, `api.airflow.us-east-1.amazonaws.com` would be `api.airflow.us-east-1.api.aws`.

Amazon MWAA uses prefixes for service endpoints, where each prefix corresponds to a set of API operations:

- Environment management – **`api`**`.airflow.`*`region`*`.amazonaws.com` (use **`api`**`.airflow.`*`region`*`.api.aws` for IPv6)
  - CreateEnvironment
  - DeleteEnvironment
  - GetEnvironment
  - ListEnvironments
  - ListTagsForResource
  - TagResource
  - UntagResource
  - UpdateEnvironment
- Airflow environment – **`env`**`.airflow.`*`region`*`.amazonaws.com` (use **`env`**`.airflow.`*`region`*`.api.aws` for IPv6)
  - CreateCliToken
  - CreateWebLoginToken

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Environments per account per Region (small, medium, large) | Each supported Region: 10 | Yes | The maximum number of environments per account per Region. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Larger environments per account per Region (xlarge, 2xlarge) | Each supported Region: 5 | Yes | The maximum number of mw1.xlarge and mw1.2xlarge environments per account per Region. |
| Web servers per environment | Each supported Region: 5 | Yes | The maximum number of web servers per environment. |
| Workers per environment | Each supported Region: 25 | Yes | The maximum number of workers per environment. |

# AWS Marketplace endpoints and quotas

AWS Marketplace is a curated digital catalog that makes it easy for customers to find, buy, deploy, and manage third-party software and services that customers need to build solutions and run their businesses.

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The AWS Marketplace website is available globally. The AWS Marketplace console is available in the US East (N. Virginia) Region. The product vendor determines the Regions in which their products are available.

## AWS Marketplace Catalog API

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | catalog.marketplace.us-east-1.amazonaws.com<br><br>catalog-marketplace.us-east-1.api.aws | HTTPS<br><br>HTTPS | |

## AWS Marketplace Commerce Analytics

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | marketplacecommerceanalytics.us-east-1.amazonaws.com | HTTPS | |

## AWS Marketplace Entitlement Service

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | entitlement.marketplace.us-east-1.amazonaws.com<br><br>entitlement-marketplace.us-east-1.api.aws | HTTPS<br><br>HTTPS | |

## AWS Marketplace Metering Service

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | metering.marketplace.us-east-2.amazonaws.com<br>metering-marketplace.us-east-2.api.aws | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | metering.marketplace.us-east-1.amazonaws.com<br>metering-marketplace.us-east-1.api.aws | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | metering.marketplace.us-west-1.amazonaws.com<br>metering-marketplace.us-west-1.api.aws | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | metering.marketplace.us-west-2.amazonaws.com<br>metering-marketplace.us-west-2.api.aws | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | metering.marketplace.af-south-1.amazonaws.com<br>metering-marketplace.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | metering.marketplace.ap-east-1.amazonaws.com<br>metering-marketplace.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | metering.marketplace.ap-south-2.amazonaws.com<br>metering-marketplace.ap-south-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | metering.marketplace.ap-southeast-3.amazonaws.com<br><br>metering-marketplace.ap-southeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | metering.marketplace.ap-southeast-4.amazonaws.com<br><br>metering-marketplace.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | metering.marketplace.ap-south-1.amazonaws.com<br><br>metering-marketplace.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | metering.marketplace.ap-northeast-3.amazonaws.com<br><br>metering-marketplace.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | metering.marketplace.ap-northeast-2.amazonaws.com<br><br>metering-marketplace.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | metering.marketplace.ap-southeast-1.amazonaws.com<br><br>metering-marketplace.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | metering.marketplace.ap-southeast-2.amazonaws.com<br><br>metering-marketplace.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | metering.marketplace.ap-northeast-1.amazonaws.com<br><br>metering-marketplace.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | metering.marketplace.ca-central-1.amazonaws.com<br><br>metering-marketplace.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | metering.marketplace.eu-central-1.amazonaws.com<br><br>metering-marketplace.eu-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | metering.marketplace.eu-west-1.amazonaws.com<br><br>metering-marketplace.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (London) | eu-west-2 | metering.marketplace.eu-west-2.amazonaws.com<br><br>metering-marketplace.eu-west-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Milan) | eu-south-1 | metering.marketplace.eu-south-1.amazonaws.com<br><br>metering-marketplace.eu-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Paris) | eu-west-3 | metering.marketplace.eu-west-3.amazonaws.com<br><br>metering-marketplace.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | metering.marketplace.eu-south-2.amazonaws.com<br>metering-marketplace.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | metering.marketplace.eu-north-1.amazonaws.com<br>metering-marketplace.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | metering.marketplace.eu-central-2.amazonaws.com<br>metering-marketplace.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | metering.marketplace.il-central-1.amazonaws.com<br>metering-marketplace.il-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | metering.marketplace.me-south-1.amazonaws.com<br>metering-marketplace.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | metering.marketplace.me-central-1.amazonaws.com<br>metering-marketplace.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | metering.marketplace.sa-east-1.amazonaws.com<br>metering-marketplace.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | metering.marketplace.us-gov-east-1.amazonaws.com<br><br>metering-marketplace.us-gov-east-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | metering.marketplace.us-gov-west-1.amazonaws.com<br><br>metering-marketplace.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

### Service quotas

For more information, see Quotas for the AWS Marketplace API.

## Amazon Mechanical Turk endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

### Service endpoints

| Region | Endpoint | Protocol |
|---|---|---|
| Sandbox endpoint for Amazon Mechanical Turk actions. | mturk-requester-sandbox.us-east-1.amazonaws.com | HTTPS |

| Region | Endpoint | Protocol |
| --- | --- | --- |
| Production endpoint for Amazon Mechanical Turk actions. | mturk-requester.us-east-1.amazonaws.com | HTTPS |

### Service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Monthly Usage | Each supported Region: 0.02 | Yes | The maximum monthly spend in USD |

## Amazon Managed Streaming for Apache Kafka endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

### Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | kafka.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | kafka-api-fips.us-east-2.api.aws | HTTPS | |
| | | kafka-fips.us-east-2.amazonaws.com | HTTPS | |
| | | kafka-api.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | kafka.us-east-1.amazonaws.com | HTTPS | |
| | | kafka-api-fips.us-east-1.api.aws | HTTPS | |
| | | kafka-fips.us-east-1.amazonaws.com | HTTPS | |
| | | kafka-api.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | kafka.us-west-1.amazonaws.com | HTTPS | |
| | | kafka-api-fips.us-west-1.api.aws | HTTPS | |
| | | kafka-fips.us-west-1.amazonaws.com | HTTPS | |
| | | kafka-api.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | kafka.us-west-2.amazonaws.com | HTTPS | |
| | | kafka-api-fips.us-west-2.api.aws | HTTPS | |
| | | kafka-fips.us-west-2.amazonaws.com | HTTPS | |
| | | kafka-api.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | kafka.af-south-1.amazonaws.com | HTTPS | |
| | | kafka-api.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | kafka.ap-east-1.amazonaws.com | HTTPS | |
| | | kafka-api.ap-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | kafka.ap-south-2.amazonaws.com<br>kafka-api.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | kafka.ap-southeast-3.amazonaws.com<br>kafka-api.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | kafka.ap-southeast-5.amazonaws.com<br>kafka-api.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | kafka.ap-southeast-4.amazonaws.com<br>kafka-api.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | kafka.ap-south-1.amazonaws.com<br>kafka-api.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | kafka.ap-southeast-6.amazonaws.com<br>kafka-api.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | kafka.ap-northeast-3.amazonaws.com<br>kafka-api.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | kafka.ap-northeast-2.amazonaws.com<br>kafka-api.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | kafka.ap-southeast-1.amazonaws.com | HTTPS |
| | | kafka-api.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | kafka.ap-southeast-2.amazonaws.com | HTTPS |
| | | kafka-api.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | kafka.ap-east-2.amazonaws.com | HTTPS |
| | | kafka-api.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | kafka.ap-southeast-7.amazonaws.com | HTTPS |
| | | kafka-api.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | kafka.ap-northeast-1.amazonaws.com | HTTPS |
| | | kafka-api.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | kafka.ca-central-1.amazonaws.com | HTTPS |
| | | kafka-api-fips.ca-central-1.api.aws | HTTPS |
| | | kafka-fips.ca-central-1.amazonaws.com | HTTPS |
| | | kafka-api.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | kafka.ca-west-1.amazonaws.com | HTTPS |
| | | kafka-api-fips.ca-west-1.api.aws | HTTPS |
| | | kafka-fips.ca-west-1.amazonaws.com | HTTPS |
| | | kafka-api.ca-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | kafka.eu-central-1.amazonaws.com<br>kafka-api.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | kafka.eu-west-1.amazonaws.com<br>kafka-api.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | kafka.eu-west-2.amazonaws.com<br>kafka-api.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | kafka.eu-south-1.amazonaws.com<br>kafka-api.eu-south-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | kafka.eu-west-3.amazonaws.com<br>kafka-api.eu-west-3.api.aws | HTTPS<br>HTTPS | |
| Europe (Spain) | eu-south-2 | kafka.eu-south-2.amazonaws.com<br>kafka-api.eu-south-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | kafka.eu-north-1.amazonaws.com<br>kafka-api.eu-north-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Zurich) | eu-central-2 | kafka.eu-central-2.amazonaws.com<br>kafka-api.eu-central-2.api.aws | HTTPS<br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | kafka.il-central-1.amazonaws.com<br>kafka-api.il-central-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | kafka.mx-central-1.amazonaws.com<br>kafka-api.mx-central-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | kafka.me-south-1.amazonaws.com<br>kafka-api.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | kafka.me-central-1.amazonaws.com<br>kafka-api.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | kafka.sa-east-1.amazonaws.com<br>kafka-api.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | kafka.us-gov-east-1.amazonaws.com<br>kafka-api.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | kafka.us-gov-west-1.amazonaws.com<br>kafka-api.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of brokers per KRaft cluster | Each supported Region: 60 | Yes | The maximum number of brokers that a KRaft cluster can contain. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Number of brokers per account | Each supported Region: 90 | Yes | The maximum number of brokers that can be created per account. |
| Number of brokers per cluster | Each supported Region: 30 | Yes | The maximum number of brokers that a cluster can contain. |
| Number of configurations per account | Each supported Region: 100 | Yes | The maximum number of custom configurations that can be created per account. |
| Number of data channels per cluster | Each supported Region: 50 | Yes | The maximum number of data channels that can be created per cluster. |
| Number of replicators per account | Each supported Region: 15 | Yes | The maximum number of replicators that can be created per account. |
| Number of revisions per configuration | Each supported Region: 50 | Yes | The maximum number of revisions that can be made to a custom configuration. |

## Amazon MSK Connect endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | kafkaconnect.us-east-2.amazonaws.com | HTTPS | |
| | | kafkaconnect-fips.us-east-2.api.aws | HTTPS | |
| | | kafkaconnect-fips.us-east-2.amazonaws.com | HTTPS | |
| | | kafkaconnect.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | kafkaconnect.us-east-1.amazonaws.com | HTTPS | |
| | | kafkaconnect-fips.us-east-1.api.aws | HTTPS | |
| | | kafkaconnect-fips.us-east-1.amazonaws.com | HTTPS | |
| | | kafkaconnect.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | kafkaconnect.us-west-1.amazonaws.com | HTTPS | |
| | | kafkaconnect-fips.us-west-1.api.aws | HTTPS | |
| | | kafkaconnect-fips.us-west-1.amazonaws.com | HTTPS | |
| | | kafkaconnect.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | kafkaconnect.us-west-2.amazonaws.com | HTTPS | |
| | | kafkaconnect-fips.us-west-2.api.aws | HTTPS | |
| | | kafkaconnect-fips.us-west-2.amazonaws.com | HTTPS | |
| | | kafkaconnect.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | kafkaconnect.af-south-1.amazonaws.com | HTTPS | |
| | | kafkaconnect.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | kafkaconnect.ap-east-1.amazonaws.com<br>kafkaconnect.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | kafkaconnect.ap-south-2.amazonaws.com<br>kafkaconnect.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | kafkaconnect.ap-southeast-3.amazonaws.com<br>kafkaconnect.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | kafkaconnect.ap-southeast-5.amazonaws.com<br>kafkaconnect.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | kafkaconnect.ap-southeast-4.amazonaws.com<br>kafkaconnect.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | kafkaconnect.ap-south-1.amazonaws.com<br>kafkaconnect.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | kafkaconnect.ap-southeast-6.amazonaws.com<br>kafkaconnect.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | kafkaconnect.ap-northeast-3.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-northeast-3.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | kafkaconnect.ap-northeast-2.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | kafkaconnect.ap-southeast-1.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | kafkaconnect.ap-southeast-2.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | kafkaconnect.ap-east-2.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-east-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | kafkaconnect.ap-southeast-7.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | kafkaconnect.ap-northeast-1.amazonaws.com | HTTPS | |
| | | kafkaconnect.ap-northeast-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | kafkaconnect.ca-central-1.amazonaws.com | HTTPS |
| | | kafkaconnect-fips.ca-central-1.api.aws | HTTPS |
| | | kafkaconnect-fips.ca-central-1.amazonaws.com | HTTPS |
| | | kafkaconnect.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | kafkaconnect.ca-west-1.amazonaws.com | HTTPS |
| | | kafkaconnect-fips.ca-west-1.api.aws | HTTPS |
| | | kafkaconnect-fips.ca-west-1.amazonaws.com | HTTPS |
| | | kafkaconnect.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | kafkaconnect.eu-central-1.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | kafkaconnect.eu-west-1.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | kafkaconnect.eu-west-2.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | kafkaconnect.eu-south-1.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | kafkaconnect.eu-west-3.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | kafkaconnect.eu-south-2.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | kafkaconnect.eu-north-1.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | kafkaconnect.eu-central-2.amazonaws.com | HTTPS |
| | | kafkaconnect.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | kafkaconnect.il-central-1.amazonaws.com | HTTPS |
| | | kafkaconnect.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | kafkaconnect.mx-central-1.amazonaws.com | HTTPS |
| | | kafkaconnect.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | kafkaconnect.me-south-1.amazonaws.com | HTTPS |
| | | kafkaconnect.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | kafkaconnect.me-central-1.amazonaws.com | HTTPS |
| | | kafkaconnect.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | kafkaconnect.sa-east-1.amazonaws.com | HTTPS |
| | | kafkaconnect.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | kafkaconnect.us-gov-east-1.amazonaws.com | HTTPS |
| | | kafkaconnect.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | kafkaconnect.us-gov-west-1.amazonaws.com<br><br>kafkaconnect.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

### Service quotas

| Dimension | Quota |
|---|---|
| Maximum custom plugins | 100 |
| Maximum worker configurations | 100 |
| Maximum connect workers | Up to 60 connect workers. If a connector is set up to have auto scaled capacity, MSK Connect uses the maximum number of workers configured for that connector to calculate the quota for the account. |

## AWS Elemental MediaConnect endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | mediaconnect.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | mediaconnect.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | mediaconnect.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | mediaconnect.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | mediaconnect.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | mediaconnect.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | mediaconnect.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | mediaconnect.ap-southeast-5.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | mediaconnect.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | mediaconnect.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | mediaconnect.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | mediaconnect.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | mediaconnect.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | mediaconnect.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | mediaconnect.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | mediaconnect.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | mediaconnect.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | mediaconnect.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | mediaconnect.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | mediaconnect.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | mediaconnect.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | mediaconnect.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | mediaconnect.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Bridges | Each supported Region: 40 | No | The maximum number of bridges that a gateway can have. |
| Entitlements | Each supported Region: 50 | No | The maximum number of entitlements that you can grant on a flow. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Flows | Each supported Region: 20 | Yes | The maximum number of flows that you can create in each AWS Region. |
| Outputs | Each supported Region: 50 | No | The maximum number of outputs that a flow can have. |
| RouterInputs | Each supported Region: 20 | Yes | The maximum number of router inputs that you can create in an AWS region |
| RouterNetworkInterfaces | Each supported Region: 10 | Yes | The maximum number of router network interfaces that you can create in an AWS region |
| RouterOutputs | Each supported Region: 20 | Yes | The maximum number of router outputs that you can create in an AWS region |

For more information, see Quotas in the *AWS Elemental MediaConnect User Guide*.

## AWS Elemental MediaConvert endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | mediaconvert.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | mediaconvert.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | mediaconvert.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | mediaconvert.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | mediaconvert.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | mediaconvert.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | mediaconvert.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | mediaconvert.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | mediaconvert.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | mediaconvert.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | mediaconvert.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | mediaconvert.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | mediaconvert.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | mediaconvert.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | mediaconvert.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | mediaconvert.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | mediaconvert.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | mediaconvert.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | mediaconvert.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | mediaconvert.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | mediaconvert.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | mediaconvert.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Concurrent jobs per account | us-east-1: 2,000<br>us-west-2: 2,000<br>eu-west-1: 2,000<br>Each of the other supported Regions: 1,000 | No | The maximum number of jobs that MediaConvert can process at one time, across all of your on-demand queues in the current AWS Region. For more information about configuring on-demand queues, see: https://docs.aws.amazon.com/ |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | mediaconvert/latest/ug/working-with-on-demand-queues.html |
| Concurrent jobs per on-demand queue | us-east-1: 500<br>us-east-2: 250<br>us-west-2: 500<br>ap-northeast-1: 250<br>ap-south-1: 250<br>eu-west-1: 500<br>eu-west-2: 250<br>Each of the other supported Regions: 100 | Yes | The maximum number of concurrent jobs that are available when you configure an on-demand queue in the current AWS Region. For more information about configuring on-demand queues, see: https://docs.aws.amazon.com/mediaconvert/latest/ug/working-with-on-demand-queues.html |
| Concurrent jobs per on-demand queue, default | us-east-1: 200<br>us-west-2: 200<br>eu-west-1: 200<br>Each of the other supported Regions: 100 | No | The default value for concurrent jobs that MediaConvert applies to new on-demand queues in the current AWS Region. For more information about configuring on-demand queues, see: https://docs.aws.amazon.com/mediaconvert/latest/ug/working-with-on-demand-queues.html |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent jobs queries per account | Each supported Region: 10 | No | The maximum number of concurrent jobs queries you can create in this account in the current AWS Region. |
| Custom job templates | Each supported Region: 100 | Yes | The maximum number of custom job templates that you can create in this account in the current AWS Region. |
| Custom output presets | Each supported Region: 100 | Yes | The maximum number of custom output presets that you can create in this account in the current AWS Region. |
| Queues per AWS Region | Each supported Region: 10 | Yes | The maximum number of on-demand and reserved queues combined that you can create in the current AWS Region. |
| Request rate for AssociateCertificate | Each supported Region: 2 per second | No | The maximum number of AssociateCertificate requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for AssociateCertificate, in a burst | Each supported Region: 100 | No | The maximum number of AssociateCertificate requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for CancelJob | Each supported Region: 2 per second | Yes | The maximum number of CancelJob requests per second that you can send in this account in the current AWS Region. |
| Request rate for CancelJob, in a burst | Each supported Region: 100 | Yes | The maximum number of CancelJob requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for CreateJob | Each supported Region: 5 per second | Yes | The maximum number of CreateJob requests per second that you can send in this account in the current AWS Region. |
| Request rate for CreateJob, in a burst | Each supported Region: 100 | Yes | The maximum number of CreateJob requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for CreateJobTemplate | Each supported Region: 2 per second | Yes | The maximum number of CreateJobTemplate requests per second that you can send in this account in the current AWS Region. |
| Request rate for CreateJobTemplate, in a burst | Each supported Region: 100 | Yes | The maximum number of CreateJobTemplate requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for CreatePreset | Each supported Region: 2 per second | Yes | The maximum number of CreatePreset requests per second that you can send in this account in the current AWS Region. |
| Request rate for CreatePreset, in a burst | Each supported Region: 100 | Yes | The maximum number of CreatePreset requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for CreateQueue | Each supported Region: 2 per second | Yes | The maximum number of CreateQueue requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for CreateQueue, in a burst | Each supported Region: 100 | Yes | The maximum number of CreateQueue requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for CreateResourceShare | Each supported Region: 0.1 | No | The maximum number of Share requests per second that you can send in this account in the current AWS Region. |
| Request rate for CreateResourceShare, in a burst | Each supported Region: 10 | No | The maximum number of Share requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for DeleteJobTemplate | Each supported Region: 2 per second | Yes | The maximum number of DeleteJobTemplate requests per second that you can send in this account in the current AWS Region. |
| Request rate for DeleteJobTemplate, in a burst | Each supported Region: 100 | Yes | The maximum number of DeleteJobTemplate requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for DeletePolicy | Each supported Region: 2 per second | No | The maximum number of DeletePolicy requests per second that you can send in this account in the current AWS Region. |
| Request rate for DeletePolicy, in a burst | Each supported Region: 100 | No | The maximum number of DeletePolicy requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for DeletePreset | Each supported Region: 2 per second | Yes | The maximum number of DeletePreset requests per second that you can send in this account in the current AWS Region. |
| Request rate for DeletePreset, in a burst | Each supported Region: 100 | Yes | The maximum number of DeletePreset requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for DeleteQueue | Each supported Region: 2 per second | Yes | The maximum number of DeleteQueue requests per second that you can send in this account in the current AWS Region. |
| Request rate for DeleteQueue, in a burst | Each supported Region: 100 | Yes | The maximum number of DeleteQueue requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for DisassociateCertificate | Each supported Region: 2 per second | No | The maximum number of DisassociateCertificate requests per second that you can send in this account in the current AWS Region. |
| Request rate for DisassociateCertificate, in a burst | Each supported Region: 100 | No | The maximum number of DisassociateCertificate requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for GetJob | Each supported Region: 2 per second | Yes | The maximum number of GetJob requests per second that you can send in this account in the current AWS Region. |
| Request rate for GetJob, in a burst | Each supported Region: 100 | Yes | The maximum number of GetJob requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for GetJobTemplate | Each supported Region: 2 per second | Yes | The maximum number of GetJobTemplate requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for GetJobTemplate, in a burst | Each supported Region: 100 | Yes | The maximum number of GetJobTemplate requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for GetJobsQueryResults | Each supported Region: 2 per second | No | The maximum number of GetJobsQueryResults requests per second that you can send in this account in the current AWS Region. |
| Request rate for GetJobsQueryResults, in a burst | Each supported Region: 100 | No | The maximum number of GetJobsQueryResults requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for GetPolicy | Each supported Region: 2 per second | No | The maximum number of GetPolicy requests per second that you can send in this account in the current AWS Region. |
| Request rate for GetPolicy, in a burst | Each supported Region: 100 | No | The maximum number of GetPolicy requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for GetPreset | Each supported Region: 2 per second | Yes | The maximum number of GetPreset requests per second that you can send in this account in the current AWS Region. |
| Request rate for GetPreset, in a burst | Each supported Region: 100 | Yes | The maximum number of GetPreset requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for GetQueue | Each supported Region: 2 per second | Yes | The maximum number of GetQueue requests per second that you can send in this account in the current AWS Region. |
| Request rate for GetQueue, in a burst | Each supported Region: 100 | Yes | The maximum number of GetQueue requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for ListJobTemplates | Each supported Region: 2 per second | Yes | The maximum number of ListJobTemplates requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for ListJobTemplates, in a burst | Each supported Region: 100 | Yes | The maximum number of ListJobTemplates requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for ListJobs | Each supported Region: 2 per second | Yes | The maximum number of ListJobs requests per second that you can send in this account in the current AWS Region. |
| Request rate for ListJobs, in a burst | Each supported Region: 100 | Yes | The maximum number of ListJobs requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for ListPresets | Each supported Region: 2 per second | Yes | The maximum number of ListPresets requests per second that you can send in this account in the current AWS Region. |
| Request rate for ListPresets, in a burst | Each supported Region: 100 | Yes | The maximum number of ListPresets requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for ListQueues | Each supported Region: 2 per second | Yes | The maximum number of ListQueues requests per second that you can send in this account in the current AWS Region. |
| Request rate for ListQueues, in a burst | Each supported Region: 100 | Yes | The maximum number of ListQueues requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for ListTagsForResource | Each supported Region: 2 per second | Yes | The maximum number of ListTagsForResource requests per second that you can send in this account in the current AWS Region. |
| Request rate for ListTagsForResource, in a burst | Each supported Region: 100 | Yes | The maximum number of ListTagsForResource requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for ListVersions | Each supported Region: 2 | Yes | The maximum number of ListVersions requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for ListVersions, in a burst | Each supported Region: 100 | Yes | The maximum number of ListVersions requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for Probe | Each supported Region: 5 | Yes | The maximum number of Probe requests per second that you can send in this account in the current AWS Region. |
| Request rate for Probe, in a burst | Each supported Region: 100 | Yes | The maximum number of Probe requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for PutPolicy | Each supported Region: 2 per second | No | The maximum number of PutPolicy requests per second that you can send in this account in the current AWS Region. |
| Request rate for PutPolicy, in a burst | Each supported Region: 100 | No | The maximum number of PutPolicy requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for SearchJobs | Each supported Region: 1 per second | No | The maximum number of SearchJobs requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for SearchJobs, in a burst | Each supported Region: 20 | No | The maximum number of SearchJobs requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for StartJobsQuery | Each supported Region: 1 per second | No | The maximum number of StartJobsQuery requests per second that you can send in this account in the current AWS Region. |
| Request rate for StartJobsQuery, in a burst | Each supported Region: 10 | No | The maximum number of StartJobsQuery requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for TagResource | Each supported Region: 2 per second | Yes | The maximum number of TagResource requests per second that you can send in this account in the current AWS Region. |
| Request rate for TagResource, in a burst | Each supported Region: 100 | Yes | The maximum number of TagResource requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for UntagResource | Each supported Region: 2 per second | Yes | The maximum number of UntagResource requests per second that you can send in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Request rate for UntagResource, in a burst | Each supported Region: 100 | Yes | The maximum number of UntagResource requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for UpdateJobTemplate | Each supported Region: 2 per second | Yes | The maximum number of UpdateJobTemplate requests per second that you can send in this account in the current AWS Region. |
| Request rate for UpdateJobTemplate, in a burst | Each supported Region: 100 | Yes | The maximum number of UpdateJobTemplate requests that you can send in one burst in this account in the current AWS Region. |
| Request rate for UpdatePreset | Each supported Region: 2 per second | Yes | The maximum number of UpdatePreset requests per second that you can send in this account in the current AWS Region. |
| Request rate for UpdatePreset, in a burst | Each supported Region: 100 | Yes | The maximum number of UpdatePreset requests that you can send in one burst in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Request rate for UpdateQueue | Each supported Region: 2 per second | Yes | The maximum number of UpdateQueue requests per second that you can send in this account in the current AWS Region. |
| Request rate for UpdateQueue, in a burst | Each supported Region: 100 | Yes | The maximum number of UpdateQueue requests that you can send in one burst in this account in the current AWS Region. |
| Reserved transcode slots (RTS) per AWS Region, per queue | Each supported Region: 30 | Yes | The maximum number of reserved transcode slots that you can add to each reserved queue in the current AWS Region. |

# AWS Elemental MediaLive endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

When you submit requests using the AWS CLI or SDKs, either leave the Region and endpoint unspecified, or specify us-east-1 as the Region. When you submit requests using the MediaLive API, use the us-east-1 Region to sign requests.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | medialive.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | medialive.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | medialive.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | medialive.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | medialive.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | medialive.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | medialive.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | medialive.ap-northeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | medialive.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | medialive.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | medialive.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | medialive.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | medialive.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | medialive.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | medialive.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | medialive.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | medialive.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | medialive.eu-north-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | medialive.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | medialive.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Anywhere Inputs | Each supported Region: 300 | Yes | The maximum number of inputs that are configured in a MediaLive Anywhere network that you can create in the current region. These inputs are push inputs but they do not count towards the standard Push Inputs quota. They count only towards this quota. |
| CDI Channels | Each supported Region: 2 | Yes | The maximum number of CDI channels (channels that include one or more CDI inputs) that you can create in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Channel Placement Groups per Cluster | Each supported Region: 100 | Yes | The maximum number of AWS MediaLive Anywhere channel placement groups that you can create in each cluster in the current region. |
| Channels | Each supported Region: 5 | Yes | The maximum number of channels that you can create in the current region. |
| CloudWatch Alarm Template Groups | Each supported Region: 20 | Yes | The maximum number of cloudwatch alarm template groups that you can create in the current region. |
| CloudWatch Alarm Templates | Each supported Region: 200 | Yes | The maximum number of cloudwatch alarm templates that you can create in the current region. |
| Clusters | Each supported Region: 100 | Yes | The maximum number of AWS MediaLive Anywhere clusters that you can create in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Device Inputs | Each supported Region: 100 | Yes | The maximum number of device inputs (an input from a device) that you can create in the current region. |
| EventBridge Rule Template Groups | Each supported Region: 20 | Yes | The maximum number of eventbridge rule template groups that you can create in the current region. |
| EventBridge Rule Templates | Each supported Region: 200 | Yes | The maximum number of eventbridge rule templates that you can create in the current region. |
| HEVC Channels | Each supported Region: 5 | Yes | The maximum number of HEVC channels (channels that include one or more HEVC outputs) that you can create in the current region. |
| Input Security Groups | Each supported Region: 5 | Yes | The maximum number of input security groups that you can create in the current region. |
| MediaConnect Inputs | Each supported Region: 100 | Yes | The maximum number of MediaConnect inputs that you can create in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| MediaLive DescribeThumbnails TPS | Each supported Region: 40 | Yes | The maximum number of transactions per second that an account can make to the DescribeThumbnails API in the current region. |
| Multiplexes | Each supported Region: 2 | Yes | The maximum number of multiplexes that you can create in the current region. |
| Networks | Each supported Region: 100 | Yes | The maximum number of AWS MediaLive Anywhere networks that you can create in the current region. |
| Nodes per Cluster | Each supported Region: 100 | Yes | The maximum number of AWS MediaLive Anywhere nodes that you can create in each cluster in the current region. |
| Pull Inputs | Each supported Region: 100 | Yes | The maximum number of pull inputs that you can create in the current region. |
| Push Inputs | Each supported Region: 5 | Yes | The maximum number of push inputs that you can create in the current region. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Reservations | Each supported Region: 50 | Yes | The maximum number of reservations that you can create in the current region. |
| SDI sources | Each supported Region: 300 | Yes | The maximum number of AWS MediaLive Anywhere SDI sources that you can create in the current region. |
| Signal Maps | Each supported Region: 30 | Yes | The maximum number of signal maps that you can create in the current region. |
| UHD Channels | Each supported Region: 1 | Yes | The maximum number of UHD channels (channels that include one or more UHD outputs) that you can create in the current region. |
| VPC Inputs | Each supported Region: 50 | Yes | The maximum number of VPC inputs that you can create in the current region. |

## AWS Elemental MediaPackage endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

These are the endpoints for MediaPackage V2 live content workflows.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | mediapackagev2.us-east-2.amazonaws.com<br>mediapackagev2-fips.us-east-2.amazonaws.com<br>mediapackagev2-fips.us-east-2.api.aws<br>mediapackagev2.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | mediapackagev2.us-east-1.amazonaws.com<br>mediapackagev2-fips.us-east-1.amazonaws.com<br>mediapackagev2-fips.us-east-1.api.aws<br>mediapackagev2.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | mediapackagev2.us-west-1.amazonaws.com<br>mediapackagev2-fips.us-west-1.amazonaws.com<br>mediapackagev2-fips.us-west-1.api.aws<br>mediapackagev2.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | mediapackagev2.us-west-2.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | mediapackagev2-fips.us-west-2.amazonaws.com<br>mediapackagev2-fips.us-west-2.api.aws<br>mediapackagev2.us-west-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | mediapackagev2.ap-south-2.amazonaws.com<br>mediapackagev2.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | mediapackagev2.ap-southeast-5.amazonaws.com<br>mediapackagev2.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | mediapackagev2.ap-southeast-4.amazonaws.com<br>mediapackagev2.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | mediapackagev2.ap-south-1.amazonaws.com<br>mediapackagev2.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | mediapackagev2.ap-northeast-3.amazonaws.com<br>mediapackagev2.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | mediapackagev2.ap-northeast-2.amazonaws.com<br>mediapackagev2.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | mediapackagev2.ap-southeast-1.amazonaws.com<br>mediapackagev2.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | mediapackagev2.ap-southeast-2.amazonaws.com<br>mediapackagev2.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | mediapackagev2.ap-northeast-1.amazonaws.com<br>mediapackagev2.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | mediapackagev2.ca-central-1.amazonaws.com<br>mediapackagev2-fips.ca-central-1.amazonaws.com<br>mediapackagev2-fips.ca-central-1.api.aws<br>mediapackagev2.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | mediapackagev2.eu-central-1.amazonaws.com<br>mediapackagev2.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | mediapackagev2.eu-west-1.amazonaws.com<br>mediapackagev2.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | mediapackagev2.eu-west-2.amazonaws.com<br>mediapackagev2.eu-west-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | mediapackagev2.eu-west-3.amazonaws.com<br>mediapackagev2.eu-west-3.api.aws | HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | mediapackagev2.eu-north-1.amazonaws.com<br>mediapackagev2.eu-north-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | mediapackagev2.me-central-1.amazonaws.com<br>mediapackagev2.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | mediapackagev2.sa-east-1.amazonaws.com<br>mediapackagev2.sa-east-1.api.aws | HTTPS<br>HTTPS | |

These are the endpoints for MediaPackage V1 live content workflows.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | mediapackage.us-east-2.amazonaws.com<br>mediapackage.us-east-2.api.aws | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | mediapackage.us-east-1.amazonaws.com<br>mediapackage.us-east-1.api.aws | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | mediapackage.us-west-1.amazonaws.com<br>mediapackage.us-west-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | mediapackage.us-west-2.amazonaws.com<br>mediapackage.us-west-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | mediapackage.ap-southeast-4.amazonaws.com<br>mediapackage.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | mediapackage.ap-south-1.amazonaws.com<br>mediapackage.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | mediapackage.ap-northeast-3.amazonaws.com<br>mediapackage.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | mediapackage.ap-northeast-2.amazonaws.com<br>mediapackage.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | mediapackage.ap-southeast-1.amazonaws.com<br>mediapackage.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | mediapackage.ap-southeast-2.amazonaws.com<br>mediapackage.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | mediapackage.ap-northeast-1.amazonaws.com<br>mediapackage.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | mediapackage.ca-central-1.amazonaws.com<br>mediapackage.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | mediapackage.eu-central-1.amazonaws.com<br>mediapackage.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | mediapackage.eu-west-1.amazonaws.com<br>mediapackage.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | mediapackage.eu-west-2.amazonaws.com<br>mediapackage.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | mediapackage.eu-west-3.amazonaws.com<br>mediapackage.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | mediapackage.eu-north-1.amazonaws.com<br>mediapackage.eu-north-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | mediapackage.sa-east-1.amazonaws.com<br>mediapackage.sa-east-1.api.aws | HTTPS<br>HTTPS |

These are the endpoints for MediaPackage V1 video on demand (VOD) content workflows.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | mediapackage-vod.us-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | mediapackage-vod.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | mediapackage-vod.us-east-1.amazonaws.com<br>mediapackage-vod.us-east-1.api.aws | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | mediapackage-vod.us-west-1.amazonaws.com<br>mediapackage-vod.us-west-1.api.aws | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | mediapackage-vod.us-west-2.amazonaws.com<br>mediapackage-vod.us-west-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | mediapackage-vod.ap-southeast-4.amazonaws.com<br>mediapackage-vod.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | mediapackage-vod.ap-south-1.amazonaws.com<br>mediapackage-vod.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | mediapackage-vod.ap-northeast-3.amazonaws.com<br>mediapackage-vod.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | mediapackage-vod.ap-northeast-2.amazonaws.com<br>mediapackage-vod.ap-northeast-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | mediapackage-vod.ap-southeast-1.amazonaws.com<br>mediapackage-vod.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | mediapackage-vod.ap-southeast-2.amazonaws.com<br>mediapackage-vod.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | mediapackage-vod.ap-northeast-1.amazonaws.com<br>mediapackage-vod.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | mediapackage-vod.ca-central-1.amazonaws.com<br>mediapackage-vod.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | mediapackage-vod.eu-central-1.amazonaws.com<br>mediapackage-vod.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | mediapackage-vod.eu-west-1.amazonaws.com<br>mediapackage-vod.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | mediapackage-vod.eu-west-2.amazonaws.com<br>mediapackage-vod.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | mediapackage-vod.eu-west-3.amazonaws.com<br>mediapackage-vod.eu-west-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | mediapackage-vod.eu-north-1.amazonaws.com<br><br>mediapackage-vod.eu-north-1.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | mediapackage-vod.sa-east-1.amazonaws.com<br><br>mediapackage-vod.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

These are the service quotes for MediaPackage V1.

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Assets per packaging group | Each supported Region: 10,000 | Yes | The maximum number of assets per packaging group that you can create. Remember that this is a per packaging group limit. If you have 1000 or fewer assets in a group, you dont need an asset limit increase, regardless of how many groups you have. |
| Burst rate of REST API requests (Live) | Each supported Region: 50 | No | The maximum number of REST API requests per second that you can burst to this account in this region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Burst rate of REST API requests (VOD) | Each supported Region: 50 | No | The maximum number of REST API requests per second that you can burst to this account in this Region. |
| Channels | Each supported Region: 30 | Yes | The maximum number of channels that you can create in this account in the current region. This limit doesnt affect endpoint limits. If you arent exceeding the endpoint limit per channel, you dont need to increase your endpoint limits when you increase your channels. |
| Concurrent harvest jobs | Each supported Region: 10 | Yes | The maximum number of harvest jobs that can be in progress at the same time. |
| Content retention | Each supported Region: 336 | No | The maximum number of hours that you can retain ingested content. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Endpoints per channel | Each supported Region: 10 | Yes | The maximum number of endpoints per channel that you can create. Remember that this is a per channel limit. If you have 10 or fewer endpoints on a channel, you dont need an endpoint limit increase, regardless of how many channels you have. |
| Ingest streams per asset | Each supported Region: 20 | No | The maximum number of streams per asset that you can ingest. |
| Ingest streams per channel | Each supported Region: 20 | No | The maximum number of streams per channel that you can ingest. |
| Live manifest length | Each supported Region: 5 | Yes | The maximum length of a live manifest, in minutes. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Packaging configurations per packaging group | Each supported Region: 10 | Yes | The maximum number of packaging configurations per packaging group that you can create. Remember that this is a per packaging group limit. If you have 10 or fewer configurations in a group, you dont need a configuration limit increase, regardless of how many groups you have. |
| Packaging groups | Each supported Region: 10 | Yes | The maximum number of packaging groups that you can create in this account in the current Region. |
| Rate of REST API requests (Live) | Each supported Region: 5 | No | The maximum number of REST API requests per second that you can send to this account in this region. |
| Rate of REST API requests (VOD) | Each supported Region: 5 | No | The maximum number of REST API requests per second that you can send to this account in this Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ingest requests per channel | Each supported Region: 50 | No | The maximum number of ingest requests per second allowed per channel. |
| Rate of manifest egress requests per asset | Each supported Region: 1,000 | No | The maximum number of manifest egress requests per second allowed per asset. |
| Rate of manifest egress requests per origin endpoint | Each supported Region: 5,000 | No | The maximum number of manifest egress requests per second allowed per origin endpoint. |
| Rate of segment egress requests per asset | Each supported Region: 600 | No | The maximum number of media segment egress requests per second allowed per asset. |
| Rate of segment egress requests per origin endpoint | Each supported Region: 300 | No | The maximum number of media segment egress requests per second allowed per origin endpoint. |
| Time-shifted manifest length | Each supported Region: 24 | No | The maximum length of a time-shifted manifest, in hours. |
| Tracks per ingest stream (Live) | Each supported Region: 10 | No | The maximum number of tracks per stream that you can ingest. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Tracks per ingest stream (VOD) | Each supported Region: 10 | No | The maximum number of tracks per stream that you can ingest. |

For more information, see Quotas in the *AWS Elemental MediaPackage User Guide*.

# AWS Elemental MediaTailor endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | api.mediatailor.us-east-2.amazonaws.com | HTTPS | |
| | | mediatailor.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | api.mediatailor.us-east-1.amazonaws.com | HTTPS | |
| | | mediatailor.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | api.mediatailor.us-west-2.amazonaws.com | HTTPS | |
| | | mediatailor.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Africa (Cape Town) | af-south-1 | api.mediatailor.af-south-1.amazonaws.com<br>mediatailor.af-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | api.mediatailor.ap-south-2.amazonaws.com<br>mediatailor.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | api.mediatailor.ap-southeast-5.amazonaws.com<br>mediatailor.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | api.mediatailor.ap-southeast-4.amazonaws.com<br>mediatailor.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | api.mediatailor.ap-south-1.amazonaws.com<br>mediatailor.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | api.mediatailor.ap-northeast-3.amazonaws.com<br>mediatailor.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | api.mediatailor.ap-northeast-2.amazonaws.com<br>mediatailor.ap-northeast-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | api.mediatailor.ap-southeast-1.amazonaws.com<br>mediatailor.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | api.mediatailor.ap-southeast-2.amazonaws.com<br>mediatailor.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.mediatailor.ap-northeast-1.amazonaws.com<br>mediatailor.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | api.mediatailor.ca-central-1.amazonaws.com<br>mediatailor.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | api.mediatailor.eu-central-1.amazonaws.com<br>mediatailor.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | api.mediatailor.eu-west-1.amazonaws.com<br>mediatailor.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | api.mediatailor.eu-west-2.amazonaws.com<br>mediatailor.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | api.mediatailor.eu-west-3.amazonaws.com<br>mediatailor.eu-west-3.api.aws | HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | api.mediatailor.eu-north-1.amazonaws.com<br>mediatailor.eu-north-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | api.mediatailor.me-central-1.amazonaws.com<br><br>mediatailor.me-central-1.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | api.mediatailor.sa-east-1.amazonaws.com<br><br>mediatailor.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Ad Insertion Requests | Each supported Region: 10,000 | Yes | Maximum requests per second when requesting manifests, segments, and tracking information when performing server side ad insertion. |
| Ad decision server (ADS) length | Each supported Region: 25,000 | No | The maximum number of characters in an ad decision server (ADS) specification. |
| Ad decision server (ADS) redirects | Each supported Region: 5 | No | The maximum depth of redirects that MediaTailor follows in VAST wrapper tags. |
| Ad decision server (ADS) timeout | Each supported Region: 3 Seconds | No | The maximum number of seconds that MediaTailor waits before timing out |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | on an open connection to an ad decision server (ADS). When a connection times out due to no response from the ADS, MediaTailor is unable to fill the ad avail with ads. |
| Channel Manifest Requests | Each supported Region: 50 | Yes | The maximum number of manifest requests per second for any single Channel Assembly channel. This is an account-level quota. |
| Channel outputs | Each supported Region: 5 | Yes | The maximum number of outputs per channel. |
| Channels per account | Each supported Region: 100 | Yes | The maximum number of channels per account. |
| Configurations | Each supported Region: 1,000 | No | The maximum number of configurations that MediaTailor allows. |
| Content origin length | Each supported Region: 512 | No | The maximum number of characters in a content origin specification. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Content origin server timeout | Each supported Region: 2 Seconds | No | The maximum number of seconds that MediaTailor waits before timing out on an open connection to the content origin server when requesting template manifests. Timeouts generate HTTP 504 (GatewayTimeoutException) response errors. |
| Live Sources | Each supported Region: 50 | Yes | The maximum number of live sources sources per source location. |
| Manifest Requests | Each supported Region: 400 | Yes | The maximum number of manifest requests per second across all channels. This is an account-level quota. |
| Manifest size | Each supported Region: 2 Megabytes | No | The maximum size, in MB, of any playback manifest, whether in input or output. To ensure that you stay under the quota, use gzip to compress your input manifests into MediaTailor. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Package configurations | Each supported Region: 5 | Yes | The maximum number of package configurations per live or VOD (video on demand) source. |
| Programs Per Channel | Each supported Region: 400 | Yes | The maximum number of programs per channel. |
| Segment Delivery Configurations Per Source | Each supported Region: 5 | Yes | The maximum number of segment delivery configurations per source location. |
| Session expiration | Each supported Region: 10 | No | The maximum amount of time (quota times the manifest duration) that MediaTailor allows a session to remain inactive before ending the session. Session activity can be a player request or an advance by the origin server. When the session expires, MediaTailor returns an HTTP 400 (Bad Request) response error. |
| Source Locations | Each supported Region: 100 | Yes | The maximum number of source locations per account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| VOD Sources | Each supported Region: 1,000 | Yes | The maximum number of video on demand (VOD) sources per source location. |

For more information, see Quotas in the *AWS Elemental MediaTailor User Guide*.

# AWS Migration Hub endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The migration tools that integrate with AWS Migration Hub send migration status to the Migration Hub in the home Region you choose. For information about choosing a home Region, see The AWS Migration Hub Home Region in the *AWS Migration Hub User Guide*.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | mgh.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | mgh.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | mgh.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | mgh.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | mgh.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | mgh.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | mgh.eu-west-2.amazonaws.com | HTTPS | |

### Service quotas

The quotas associated with AWS Migration Hub are the AWS Application Discovery Service quotas. For more information, see AWS Application Discovery Service Quotas.

## Migration Hub Orchestrator endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | migrationhub-orchestrator.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | migrationhub-orchestrator.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | migrationhub-orchestrator.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | migrationhub-orchestrator.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | migrationhub-orchestrator.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | migrationhub-orchestrator.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | migrationhub-orchestrator.eu-west-2.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum steps | 15 | No | The maximum number of steps in a step group. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum step groups | 15 | No | The maximum number of step groups in a migration workflow. |
| Maximum migration workflows | 50 | No | The maximum number of migration workflows in progress. |

# AWS Migration Hub Refactor Spaces endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | refactor-spaces.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | refactor-spaces.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (N. California) | us-west-1 | refactor-spaces.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | refactor-spaces.us-west-2.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | refactor-spaces.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | refactor-spaces.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | refactor-spaces.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | refactor-spaces.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | refactor-spaces.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | refactor-spaces.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | refactor-spaces.ca-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | refactor-spaces.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | refactor-spaces.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | refactor-spaces.eu-west-2.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | refactor-spaces.eu-west-3.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | refactor-spaces.eu-north-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | refactor-spaces.sa-east-1.amazonaws.com | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Applications | Each supported Region: 600 | Yes | The maximum number of Refactor Spaces applications that you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Environments | Each supported Region: 50 | Yes | The maximum number of Refactor Spaces environments that you can create in this account in the current Region. |
| Routes | Each supported Region: 1,000 | Yes | The maximum number of Refactor Spaces routes that you can create in this account in the current Region. |
| Services | Each supported Region: 500 | Yes | The maximum number of Refactor Spaces services that you can create in this account in the current Region. |

## Migration Hub Strategy Recommendations endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | migrationhub-strategy.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | migrationhub-strategy.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | migrationhub-strategy.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | migrationhub-strategy.ap-northeast-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | migrationhub-strategy.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | migrationhub-strategy.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | migrationhub-strategy.eu-west-2.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Active Assessment Maximum | Each supported Region: 1 | Yes | The maximum number of concurrent active assessments |
| Active Import Maximum | Each supported Region: 5 | Yes | The maximum number of concurrent active import tasks |
| Assessment Maximum | Each supported Region: 50 | Yes | The maximum number of assessments per AWS account |
| Maximum Server per Assessment | Each supported Region: 300 | Yes | The maximum number of servers per assessment |

# Amazon Monitron endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

**Important**

Amazon Monitron currently does not support programmatic access to service endpoints.

Amazon Monitron is currently supported in the following Regions:

- US East (N. Virginia): us-east-1
- Europe (Ireland): eu-west-1
- Asia Pacific (Sydney): ap-south-east-2

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Asset classes per project | Each supported Region: 25 | Yes | The maximum number of asset classes per project. |
| Assets per site | Each supported Region: 100 | Yes | The maximum number of assets per site. |
| Gateways per site | Each supported Region: 200 | Yes | The maximum number of gateways per site. |
| Positions per asset | Each supported Region: 20 | Yes | The maximum number of positions per asset. |
| Positions per custom class | Each supported Region: 500 | Yes | The maximum number of positions per custom class. |
| Projects per account | Each supported Region: 10 | Yes | The maximum number of projects that can be created for an account. |
| Sites per project | Each supported Region: 50 | Yes | The maximum number of sites per project. |
| Users per site | Each supported Region: 20 | Yes | The maximum number of users per site. |

# Amazon MQ endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | mq.us-east-2.amazonaws.com | HTTPS |
| | | mq-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | mq.us-east-1.amazonaws.com | HTTPS |
| | | mq-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | mq.us-west-1.amazonaws.com | HTTPS |
| | | mq-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | mq.us-west-2.amazonaws.com | HTTPS |
| | | mq-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | mq.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-east-1 | mq.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | mq.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | mq.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | mq.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | mq.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | mq.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | mq.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | mq.ap-northeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | mq.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | mq.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | mq.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | mq.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | mq.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | mq.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | mq.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | mq.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | mq.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | mq.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | mq.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | mq.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | mq.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | mq.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | mq.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | mq.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | mq.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | mq.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | mq.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | mq.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | mq.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | mq.us-gov-east-1.amazonaws.com<br>mq-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | mq.us-gov-west-1.amazonaws.com<br>mq-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| API burst limit | Each supported Region: 100 | No | The API burst limit, aggregated per AWS account, across all Amazon MQ APIs to maintain service bandwidth. This doesnt apply to ActiveMQ or RabbitMQ broker messaging APIs. For example, Amazon MQ doesnt throttle the sending or receiving of messages. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| API rate limit | Each supported Region: 15 | No | The API rate limit, aggregated per AWS account, across all Amazon MQ APIs to maintain service bandwidth. This doesnt apply to ActiveMQ or RabbitMQ broker messaging APIs. For example, Amazon MQ doesnt throttle the sending or receiving of messages. |
| Destinations monitored in CloudWatch (ActiveMQ) | Each supported Region: 1,000 | No | The maximum number of destinations (ActiveMQ queues, ActiveMQ topics) monitored in CloudWatch. |
| Destinations monitored in CloudWatch (RabbitMQ) | Each supported Region: 500 | No | The maximum number of destinations (RabbitMQ queues) monitored in CloudWatch. |
| Groups per user (simple auth) | Each supported Region: 20 | No | The maximum number of groups per user (simple auth). This does not apply to RabbitMQ brokers. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Job scheduler usage limit per broker backed by Amazon EBS | Each supported Region: 50 Gigabytes | No | The job scheduler usage limit (in GB) per broker backed by Amazon EBS. This does not apply to RabbitMQ brokers. |
| Number of brokers, per region | Each supported Region: 200 | Yes | The maximum number of brokers, per region. |
| Revisions per configuration | Each supported Region: 300 | No | The maximum number of revisions per configuration. This does not apply to RabbitMQ brokers. |
| Security groups per broker | Each supported Region: 5 | No | The maximum number of security groups per broker. |
| Storage capacity per larger broker | Each supported Region: 200 Gigabytes | No | The maximum storage capacity (in GB) per mq.m5.* instance type broker. |
| Storage capacity per smaller broker | Each supported Region: 20 Gigabytes | No | The maximum storage capacity (in GB) per smaller broker (mq.*.micro instance type brokers). |
| Tags per broker | Each supported Region: 50 | No | The maximum number of tags per broker. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Temporary storage capacity per larger broker | Each supported Region: 50 Gigabytes | No | The maximum temporary storage capacity (in GB) per larger broker (mq.*.*large instance type brokers). This does not apply to RabbitMQ brokers. |
| Temporary storage capacity per smaller broker | Each supported Region: 5 Gigabytes | No | The maximum temporary storage capacity (in GB) per smaller broker (mq.*.micro instance type brokers). This does not apply to RabbitMQ brokers. |
| Users per broker (simple auth) | Each supported Region: 250 | No | The maximum number of users per broker (simple auth). This does not apply to RabbitMQ brokers. |
| Wire-level connections per larger broker | Each supported Region: 2,000 | Yes | The maximum number of wire-level connections per larger broker (mq.*.*large instance type brokers). This does not apply to RabbitMQ brokers. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Wire-level connections per smaller broker | Each supported Region: 300 | No | The maximum number of wire-level connections per smaller broker (mq.*.micro instance type brokers). This does not apply to RabbitMQ brokers. |

For more information, see Quotas in Amazon MQ in the *Amazon MQ Developer Guide*.

# Amazon Neptune endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | rds.us-east-2.amazonaws.com<br><br>rds-fips.us-east-2.amazonaws.com | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | rds.us-east-1.amazonaws.com<br><br>rds-fips.us-east-1.amazonaws.com | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| US West (N. California) | us-west-1 | rds.us-west-1.amazonaws.com<br><br>rds-fips.us-west-1.amazonaws.com | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | rds.us-west-2.amazonaws.com<br><br>rds-fips.us-west-2.amazonaws.com | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Africa (Cape Town) | af-south-1 | rds.af-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | rds.ap-east-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | rds.ap-south-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rds.ap-southeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rds.ap-southeast-5.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rds.ap-southeast-4.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rds.ap-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | rds.ap-northeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | rds.ap-northeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | rds.ap-southeast-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | rds.ap-southeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | rds.ap-northeast-1.amazonaws.com | HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | rds.ca-central-1.amazonaws.com<br><br>rds-fips.ca-central-1.amazonaws.com | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Canada West (Calgary) | ca-west-1 | rds.ca-west-1.amazonaws.com<br><br>rds-fips.ca-west-1.amazonaws.com | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | rds.eu-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | rds.eu-west-1.amazonaws.com | HTTP and HTTPS | |
| Europe (London) | eu-west-2 | rds.eu-west-2.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | rds.eu-west-3.amazonaws.com | HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | rds.eu-south-2.amazonaws.com | HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | rds.eu-north-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Zurich) | eu-central-2 | rds.eu-central-2.amazonaws.com | HTTP and HTTPS | |
| Israel (Tel Aviv) | il-central-1 | rds.il-central-1.amazonaws.com | HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | rds.me-south-1.amazonaws.com | HTTP and HTTPS | |
| Middle East (UAE) | me-central-1 | rds.me-central-1.amazonaws.com | HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | rds.sa-east-1.amazonaws.com | HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | rds.us-gov-east-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | rds.us-gov-west-1.amazonaws.com | HTTP and HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Cluster endpoints per DB cluster | Each supported Region: 5 | Yes | The maximum number of cluster endpoints per DB cluster. |
| Cross-region snapshot copy requests | Each supported Region: 5 | Yes | The maximum number of cross-region snapshot copy requests. |
| DB cluster Roles | Each supported Region: 5 | Yes | The maximum number of DB cluster roles. |
| DB cluster manuals snapshots | Each supported Region: 100 | Yes | The maximum number of DB cluster manual snapshots. |
| DB cluster parameter groups | Each supported Region: 50 | Yes | The maximum number of DB cluster parameter groups. |
| DB clusters | Each supported Region: 40 | Yes | The maximum number of DB clusters. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DB instance parameter groups | Each supported Region: 50 | Yes | The maximum number of DB instance parameter groups. |
| DB instances | Each supported Region: 40 | Yes | The maximum number of DB instances. |
| DB subnet groups | Each supported Region: 50 | Yes | The maximum number of DB subnet groups. |
| Event subscriptions | Each supported Region: 20 | Yes | The maximum number of event subscriptions. |
| Read replicas per cluster | Each supported Region: 15 | No | The maximum number of read replicas per cluster. |
| Reserved DB instances | Each supported Region: 40 | Yes | The maximum number of reserved DB instances. |
| Tags per resource | Each supported Region: 50 | Yes | The maximum number of tags per resource. |

For more information, see Amazon Neptune quotas in the *Amazon Neptune User Guide*.

## AWS Network Firewall endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | network-firewall.us-east-2.amazonaws.com<br>network-firewall-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | network-firewall.us-east-1.amazonaws.com<br>network-firewall-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | network-firewall.us-west-1.amazonaws.com<br>network-firewall-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | network-firewall.us-west-2.amazonaws.com<br>network-firewall-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | network-firewall.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | network-firewall.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | network-firewall.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | network-firewall.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | network-firewall.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | network-firewall.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | network-firewall.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | network-firewall.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | network-firewall.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | network-firewall.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | network-firewall.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | network-firewall.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | network-firewall.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | network-firewall.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | network-firewall.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | network-firewall.ca-central-1.amazonaws.com<br>network-firewall-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | network-firewall.ca-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | network-firewall.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | network-firewall.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | network-firewall.eu-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | network-firewall.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | network-firewall.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | network-firewall.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | network-firewall.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | network-firewall.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | network-firewall.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | network-firewall.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | network-firewall.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | network-firewall.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | network-firewall.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | network-firewall.us-gov-east-1.amazonaws.com<br><br>network-firewall-fips.us-gov-east-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | network-firewall.us-gov-west-1.amazonaws.com<br><br>network-firewall-fips.us-gov-west-1.amazonaws.com | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CA certificates per TLS configuration | Each supported Region: 1 | No | The maximum number of certificate authority (CA) certificates for a TLS inspection configuration. CA certificates are used for outbound SSL/TLS inspection. |
| Container associations | Each supported Region: 100 | Yes | The maximum number of container associations per account per Region. |
| Firewall policies | Each supported Region: 20 | Yes | The maximum number of firewall policies per account per Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Firewalls | Each supported Region: 5 | Yes | The maximum number of firewalls per account per Region. |
| IP set references per Suricata compatible stateful rule group | Each supported Region: 5 | No | The maximum number of IP set references per Suricata compatible stateful rule group. |
| Network traffic bandwidth per firewall endpoint | Each supported Region: 100 | No | The maximum network traffic bandwidth, in Gbps, for any firewall endpoint. If you require more traffic bandwidth, you can split your resources into subnets and create a firewall in each subnet. |
| Number of firewalls that can use the same policy | Each supported Region: 1,000 | No | The maximum number of firewalls that can use the same firewall policy. |
| Number of policies that can use the same rule group | Each supported Region: 1,000 | No | The maximum number of firewall policies that can use the same rule group. |
| Number of policies using a TLS inspection configuration | Each supported Region: 1,000 | No | The maximum number of firewall policies that can use the same TLS inspection configuration. |
| Required firewall policies per firewall | Each supported Region: 1 | No | The required number of firewall policies per firewall. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Resource filters | Each supported Region: 50 | Yes | The maximum number of resource filters per account per Region. |
| Server certificates per TLS configuration | Each supported Region: 10 | No | The maximum number of server certificates for a TLS inspection configuration. Server certificates are used for inbound SSL/TLS inspection. |
| Stateful rule group capacity | Each supported Region: 30,000 | No | The maximum stateful rule group capacity. |
| Stateful rule groups per policy | Each supported Region: 20 | No | The maximum number of stateful rule groups per firewall policy. |
| Stateful rulegroups | Each supported Region: 50 | Yes | The maximum number of stateful rule groups per account per Region. |
| Stateful rules per policy | Each supported Region: 30,000 | Yes | The maximum number of stateful rules per firewall policy. This is the total across all rule groups that are referenced by the policy. You can increase this up to 50,000, but higher settings might impact firewall performance. |
| Stateless rule group capacity | Each supported Region: 30,000 | No | The maximum stateless rule group capacity. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Stateless rule group custom actions | Each supported Region: 10 | No | The maximum number of custom actions per stateless rule group. |
| Stateless rule groups per policy | Each supported Region: 20 | No | The maximum number of stateless rule groups per firewall policy. |
| Stateless rulegroups | Each supported Region: 50 | Yes | The maximum number of stateless rule groups per account per Region. |
| Stateless rules per policy | Each supported Region: 30,000 | No | The maximum number of stateless rules per firewall policy. This is the total across all rule groups that are referenced by the policy. |
| Suricata rule character length | Each supported Region: 8,192 | No | The maximum character length of a Suricata rule. Each variable value in the rule counts toward this limit. |
| Suricata rules string size | Each supported Region: 2,000,000 | No | The maximum size of a Suricata-compatible rules string for a rule group, in bytes. |
| TLS configurations | Each supported Region: 20 | Yes | The maximum number of TLS configurations per account per Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| TLS inspection configurations per policy | Each supported Region: 1 | No | The maximum number of TLS inspection configurations per policy. |
| VPC endpoint associations | Each supported Region: 300 | Yes | The maximum number of VPC endpoint associations per account per Region. |
| VPC endpoint associations per Availability Zone per Firewall | Each supported Region: 50 | No | The maximum number of VPC endpoint associations for a firewall per Availability Zone. |

For more information, see AWS Network Firewall quotas in the *Network Firewall Developer Guide*.

# Amazon OpenSearch Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### OpenSearch Service API

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | es.us-east-2.amazonaws.com | HTTPS |
| | | es-fips.us-east-2.amazonaws.com | HTTPS |
| | | aos.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | es.us-east-1.amazonaws.com | HTTPS |
| | | es-fips.us-east-1.amazonaws.com | HTTPS |
| | | aos.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | es.us-west-1.amazonaws.com | HTTPS |
| | | es-fips.us-west-1.amazonaws.com | HTTPS |
| | | aos.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | es.us-west-2.amazonaws.com | HTTPS |
| | | es-fips.us-west-2.amazonaws.com | HTTPS |
| | | aos.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | es.af-south-1.amazonaws.com | HTTPS |
| | | aos.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | es.ap-east-1.amazonaws.com | HTTPS |
| | | aos.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | es.ap-south-2.amazonaws.com | HTTPS |
| | | aos.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | es.ap-southeast-3.amazonaws.com<br>aos.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | es.ap-southeast-5.amazonaws.com<br>aos.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | es.ap-southeast-4.amazonaws.com<br>aos.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | es.ap-south-1.amazonaws.com<br>aos.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | es.ap-southeast-6.amazonaws.com<br>aos.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | es.ap-northeast-3.amazonaws.com<br>aos.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | es.ap-northeast-2.amazonaws.com<br>aos.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | es.ap-southeast-1.amazonaws.com<br>aos.ap-southeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | es.ap-southeast-2.amazonaws.com<br>aos.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | es.ap-east-2.amazonaws.com<br>aos.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | es.ap-southeast-7.amazonaws.com<br>aos.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | es.ap-northeast-1.amazonaws.com<br>aos.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | es.ca-central-1.amazonaws.com<br>aos.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | es.ca-west-1.amazonaws.com<br>aos.ca-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | es.eu-central-1.amazonaws.com<br>aos.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | es.eu-west-1.amazonaws.com<br>aos.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | es.eu-west-2.amazonaws.com<br>aos.eu-west-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | es.eu-south-1.amazonaws.com | HTTPS | |
| | | aos.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | es.eu-west-3.amazonaws.com | HTTPS | |
| | | aos.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | es.eu-south-2.amazonaws.com | HTTPS | |
| | | aos.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | es.eu-north-1.amazonaws.com | HTTPS | |
| | | aos.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | es.eu-central-2.amazonaws.com | HTTPS | |
| | | aos.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | es.il-central-1.amazonaws.com | HTTPS | |
| | | aos.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | es.mx-central-1.amazonaws.com | HTTPS | |
| | | aos.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | es.me-south-1.amazonaws.com | HTTPS | |
| | | aos.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | es.me-central-1.amazonaws.com | HTTPS | |
| | | aos.me-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| South America (São Paulo) | sa-east-1 | es.sa-east-1.amazonaws.com<br>aos.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | es.us-gov-east-1.amazonaws.com<br>es-fips.us-gov-east-1.amazonaws.com<br>aos.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | es.us-gov-west-1.amazonaws.com<br>es-fips.us-gov-west-1.amazonaws.com<br>aos.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

**OpenSearch Serverless API**

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | aoss.us-east-2.amazonaws.com<br>dashboards.us-east-2.aoss-fips.amazonaws.com<br>us-east-2.aoss.amazonaws.com<br>dashboards.us-east-2.aoss.amazonaws.com<br>us-east-2.aoss-fips.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | aoss.us-east-1.amazonaws.com<br>us-east-1.aoss.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | dashboards.us-east-1.aoss.amazonaws.com | HTTPS | |
| | | dashboards.us-east-1.aoss-fips.amazonaws.com | HTTPS | |
| | | us-east-1.aoss-fips.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | aoss.us-west-1.amazonaws.com | HTTPS | |
| | | dashboards.us-west-1.aoss-fips.amazonaws.com | HTTPS | |
| | | dashboards.us-west-1.aoss.amazonaws.com | HTTPS | |
| | | us-west-1.aoss.amazonaws.com | HTTPS | |
| | | us-west-1.aoss-fips.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | aoss.us-west-2.amazonaws.com | HTTPS | |
| | | dashboards.us-west-2.aoss-fips.amazonaws.com | HTTPS | |
| | | dashboards.us-west-2.aoss.amazonaws.com | HTTPS | |
| | | us-west-2.aoss.amazonaws.com | HTTPS | |
| | | us-west-2.aoss-fips.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | aoss.ap-east-1.amazonaws.com | HTTPS | |
| | | ap-east-1.aoss.amazonaws.com | HTTPS | |
| | | dashboards.ap-east-1.aoss.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | aoss.ap-south-2.amazonaws.com | HTTPS |
| | | dashboards.ap-south-2.aoss.amazonaws.com | HTTPS |
| | | ap-south-2.aoss.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | aoss.ap-south-1.amazonaws.com | HTTPS |
| | | dashboards.ap-south-1.aoss.amazonaws.com | HTTPS |
| | | ap-south-1.aoss.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | aoss.ap-northeast-3.amazonaws.com | HTTPS |
| | | ap-northeast-3.aoss.amazonaws.com | HTTPS |
| | | dashboards.ap-northeast-3.aoss.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | aoss.ap-northeast-2.amazonaws.com | HTTPS |
| | | dashboards.ap-northeast-2.aoss.amazonaws.com | HTTPS |
| | | ap-northeast-2.aoss.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | aoss.ap-southeast-1.amazonaws.com | HTTPS |
| | | dashboards.ap-southeast-1.aoss.amazonaws.com | HTTPS |
| | | ap-southeast-1.aoss.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | aoss.ap-southeast-2.amazonaws.com | HTTPS |
| | | dashboards.ap-southeast-2.aoss.amazonaws.com | HTTPS |
| | | ap-southeast-2.aoss.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | aoss.ap-northeast-1.amazonaws.com | HTTPS |
| | | dashboards.ap-northeast-1.aoss.amazonaws.com | HTTPS |
| | | ap-northeast-1.aoss.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | aoss.ca-central-1.amazonaws.com | HTTPS |
| | | ca-central-1.aoss-fips.amazonaws.com | HTTPS |
| | | dashboards.ca-central-1.aoss-fips.amazonaws.com | HTTPS |
| | | dashboards.ca-central-1.aoss.amazonaws.com | HTTPS |
| | | ca-central-1.aoss.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | aoss.eu-central-1.amazonaws.com | HTTPS |
| | | dashboards.eu-central-1.aoss.amazonaws.com | HTTPS |
| | | eu-central-1.aoss.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | aoss.eu-west-1.amazonaws.com | HTTPS |
| | | eu-west-1.aoss.amazonaws.com | HTTPS |
| | | dashboards.eu-west-1.aoss.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | aoss.eu-west-2.amazonaws.com | HTTPS |
| | | eu-west-2.aoss.amazonaws.com | HTTPS |
| | | dashboards.eu-west-2.aoss.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Milan) | eu-south-1 | aoss.eu-south-1.amazonaws.com | HTTPS |
| | | eu-south-1.aoss.amazonaws.com | HTTPS |
| | | dashboards.eu-south-1.aoss.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | aoss.eu-west-3.amazonaws.com | HTTPS |
| | | eu-west-3.aoss.amazonaws.com | HTTPS |
| | | dashboards.eu-west-3.aoss.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | aoss.eu-south-2.amazonaws.com | HTTPS |
| | | dashboards.eu-south-2.aoss.amazonaws.com | HTTPS |
| | | eu-south-2.aoss.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | aoss.eu-north-1.amazonaws.com | HTTPS |
| | | eu-north-1.aoss.amazonaws.com | HTTPS |
| | | dashboards.eu-north-1.aoss.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | aoss.eu-central-2.amazonaws.com | HTTPS |
| | | eu-central-2.aoss.amazonaws.com | HTTPS |
| | | dashboards.eu-central-2.aoss.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | aoss.sa-east-1.amazonaws.com | HTTPS |
| | | sa-east-1.aoss.amazonaws.com | HTTPS |
| | | dashboards.sa-east-1.aoss.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | aoss.us-gov-east-1.amazonaws.com | HTTPS |
| | | us-gov-east-1.aoss-fips.amazonaws.com | HTTPS |
| | | dashboards.us-gov-east-1.aoss-fips.amazonaws.com | HTTPS |
| | | us-gov-east-1.aoss.amazonaws.com | HTTPS |
| | | dashboards.us-gov-east-1.aoss.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | aoss.us-gov-west-1.amazonaws.com | HTTPS |
| | | dashboards.us-gov-west-1.aoss.amazonaws.com | HTTPS |
| | | us-gov-west-1.aoss-fips.amazonaws.com | HTTPS |
| | | us-gov-west-1.aoss.amazonaws.com | HTTPS |
| | | dashboards.us-gov-west-1.aoss-fips.amazonaws.com | HTTPS |

### OpenSearch Ingestion API

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | osis.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | osis.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (N. California) | us-west-1 | osis.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | osis.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | osis.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | osis.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | osis.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | osis.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | osis.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | osis.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | osis.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | osis.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | osis.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | osis.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | osis.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | osis.eu-north-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | osis.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | osis.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | osis.us-gov-west-1.amazonaws.com | HTTPS | |

**OpenSearch UI**

## Service quotas

For more information, see Amazon OpenSearch Service quotas.

## OpenSearch Service domain and instance quotas

Your AWS account has the following quotas related to OpenSearch Service domains:

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Dedicated master instances per domain | Each supported Region: 5 | No | The maximum number of dedicated master instances in a single Amazon OpenSearch Service domain. |
| Domains per Region | Each supported Region: 500 | Yes | The maximum number of Amazon OpenSearch Service domains you can create in each AWS Region. |
| Instances per domain | Each supported Region: 200 | Yes | The maximum number of instances in a single Amazon OpenSearch Service domain. You can request an increase up to 1002 instances per domain. |
| Instances per domain (T2 instance type) | Each supported Region: 10 | No | The maximum number of T2 instances in a single Amazon OpenSearch Service domain. |
| Warm instances per domain | Each supported Region: 150 | No | The maximum number of warm nodes in a single Amazon OpenSearch Service cluster. |

Your AWS account has the following additional OpenSearch Service limits:

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Total storage per domain | 3 PiB | No | This maximum is the sum of all data nodes and warm nodes. For example, your domain might have 45 `r6gd.16xlarge.search` instances and 140 `ultrawarm1.large.search` instances for a total of 2.88 PiB of storage. |
| Custom packages per Region | 25 | No | |
| Custom packages per domain | 20 | No | |

## OpenSearch Serverless quotas

Your AWS account has the following quotas related to OpenSearch Serverless resources.

To view the quotas for OpenSearch Serverless, open the Service Quotas console. In the navigation pane, choose **AWS services** and select **Amazon OpenSearch Serverless**.

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Default indexing capacity (OCUs) | 10 | No | |

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Default search capacity (OCUs) | 10 | No | |
| Maximum indexing capacity (OCUs) | 1,700 | No | |
| Maximum search capacity (OCUs) | 1,700 | No | |
| Data access policies | 500 | No | |
| Encryption policies | 150 | No | |
| Network policies | 500 | No | |
| SAML providers | 50 | No | |
| Data access policy size | 10,240 bytes | No | |
| Network policy size | 10,240 bytes | No | |
| SAML provider size | 51,200 bytes | No | |
| Encryption policy size | 10,240 bytes | No | |

Your AWS account has the following additional OpenSearch Serverless limits:

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Collection tags | 50 | No | |

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Collections per collection group (Classic) | 1,500 | No | The maximum number of collections per Classic collection group. The actual number of collections can be further limited by MaxOCUs based on the size of collections. |
| Collections per collection group (NextGen) | 10,000 | Yes | The maximum number of collections per NextGen collection group. You can request an increase through the Service Quotas console. |

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Collections with unique KMS keys (not in a collection group) | Maximum OCU/2 | No | This limit only applies to collections that are not part of a collection group. It depends on the maximum number of OCUs you set and how many OCUs are already in use. Collections within a collection group have a separate limit of 1,500 collections per group, regardless of the number of unique KMS keys used. |
| OpenSearch Serverless-managed VPC endpoints | 50 | No | This limit only applies to OpenSearch Serverless-managed VPC endpoints. It doesn't include OpenSearch Service-managed VPC endpoints. |

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Indexes per Amazon OpenSearch Serverless collection | 1000 | No | The maximum number of indexes per Amazon OpenSearch Serverless collection is 1000. |
| Index templates per Amazon OpenSearch Serverless collection | 500 | No | The maximum number of index templates per Amazon OpenSearch Serverless collection is 500. |
| Maximum collection groups | 300 | No | The maximum number of collection groups per account. |

## OpenSearch Ingestion quotas

Your AWS account has the following quotas related to OpenSearch Ingestion resources. Unless otherwise noted, each quota is Region-specific.

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Connected VPCs per pipeline | Each supported Region: 1 | No | |
| Pipelines connected to a single VPC domain | Each supported Region: 50 | No | |

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Pipelines | Each supported Region: 50 | No | |
| Characters per pipeline configuration | Each supported Region: 24000 | No | |
| Sources per pipeline | Each supported Region: 100000 | No | This limit only applies to *plugin* sources, such as HTTP, OTel, or S3. You can still chain sub-pipelines together within a single OpenSearch Ingestion pipeline configuration. |
| Unique OpenSearch sinks per pipeline | Each supported Region: 1 | No | This limit only applies to OpenSearch *domain* and *collection* sinks. You can still chain sub-pipelines together within a single OpenSearch Ingestion pipeline configuration. |

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| Total OpenSearch sinks per pipeline | Each supported Region: 50 | No | This limit refers to the number of times that you can specify the same OpenSearch sink within a pipeline configuration file. The number of non "pipeline " sinks within the `pipelineConfigurationBody` must be between 1 and 20, inclusive of OpenSearch Ingestion pipelines. |
| Minimum Ingestion OCUs per pipeline | Each supported Region: 1 | No | |
| Maximum Ingestion OCUs per stateless pipeline | Each supported Region: 96 | No | See Stateless versus stateful processors. |
| Maximum Ingestion OCUs per stateful pipeline | Each supported Region: 48 | No | See Stateless versus stateful processors. |

## OpenSearch UI quotas

Your AWS account has the following quotas related to OpenSearch UI. Unless otherwise noted, each quota is Region-specific.

| Name | Default | Adjustable | Notes |
| --- | --- | --- | --- |
| OpenSearch UI (Dashboards) applications per account | Each supported Region: 30 | Yes | You can request an increase up to 50 applications per account. |

# Oracle Database@AWS endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Oracle Database@AWS

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | odb.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | odb.us-west-2.amazonaws.com | HTTPS | |

# AWS Organizations endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Because AWS Organizations is a global service, there is a single global endpoint for all of the AWS Regions in each partition.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | organizations.us-east-1.amazonaws.com | HTTPS |
| | | organizations-fips.us-east-1.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | organizations.us-east-1.amazonaws.com | HTTPS |
| | | organizations-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | organizations.us-east-1.amazonaws.com | HTTPS |
| | | organizations-fips.us-east-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | organizations.us-east-1.amazonaws.com | HTTPS |
| | | organizations-fips.us-east-1.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | organizations.us-east-1.amazonaws.com | HTTPS |
| | | organizations-fips.us-east-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-east-1 | organizations.us-east-1.amazonaws.com | HTTPS |
| | | organizations-fips.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | organizations.us-east-1.amazonaws.com<br>organizations-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | organizations.us-gov-west-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | organizations.us-gov-west-1.amazonaws.com | HTTPS |

## Service quotas

AWS Organizations has the following quotas:

> **Note**
>
> AWS Organizations limits apply at the organization level, unless otherwise specified. Many quotas apply only to actions performed from the AWS Organizations management account.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Enable all features request expiration | Each supported Region: 90 | No | Maximum number of days to allow before request to enable all features expires. |
| Handshake expiration | Each supported Region: 30 | No | Maximum number of days to allow to complete handshake. |
| Inbound billing transfers | Each supported Region: 0 | Yes | The maximum number of external organizations that you can manage and pay for at any given time. The maximum value is 1,000. |
| Invitation acceptance expiration | Each supported Region: 15 | No | Maximum number of days to allow before invitations to join an organization expire. |
| Length of bill-transfer chain | Each supported Region: 1 | Yes | The maximum number of interconnections that each of your inbound billing transfers can have at any given time. The maximum value for this quota is 2 for Distributors, and 1 for Solution Providers and Direct Customers. The maximum value for Solution Providers and |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | Direct Customers cannot be increased. |
| Maximum number of accounts | Each supported Region: 10 | Yes | The maximum number of accounts allowed in an organization. |
| Member accounts you can concurrently close | Each supported Region: 3 | No | Maximum number of accounts that you can concurrently close. As soon as one finishes, you can close another. |
| Member accounts you can concurrently create | Each supported Region: 5 | No | Maximum number of accounts that you can concurrently create. As soon as one finishes, you can start another. |
| Minimum age for removal of created accounts | Each supported Region: 4 | No | The minimum number of days a created account must exist before you can remove it from the organization. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Number of accounts you can close within a 30 day period | Each supported Region: 250 | No | The maximum number of member accounts that can be closed in a 30-day period. For organizations with up to 250 accounts, the quota equals the number of accounts. For larger organizations, the quota is 250 or 20-percent of total member accounts, whichever is greater, up to 1000. |
| Number of invitation attempts you can perform in a 24-hour period | Each supported Region: 20 | No | Maximum number of invitation attempts you can perform in a 24-hour period. The number of attempts can be either 20 or the maximum number of accounts allowed in your organization, whichever is greater. |
| OU maximum nesting in a root | Each supported Region: 5 | No | The maximum number of levels that you can nest OUs under the root. |
| OUs in an organization | Each supported Region: 2,000 | No | The maximum number of OUs that you can include in an organization. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Outbound billing transfers | Each supported Region: 1 | No | The maximum number of external organizations that you can allow to manage and pay for your consolidated bill at any given time. |
| Roots in an organization | Each supported Region: 1 | No | The maximum number of root accounts allowed per organization. |
| Service control policies in an organization | Each supported Region: 10,000 | No | The maximum number of service control policies (SCPs) allowed per organization. |
| Service control policies per OU | Each supported Region: 10 | No | The maximum number of service control policies (SCPs) allowed per OU. |
| Service control policies per account | Each supported Region: 10 | No | The maximum number of service control policies (SCPs) allowed per account. |
| Service control policies per root | Each supported Region: 10 | No | The maximum number of service control policies (SCPs) allowed per root. |
| Service control policy document size | Each supported Region: 10,240 | No | The maximum document size (in characters) allowed for service control policies (SCPs). |

For more information, see Quotas and service limits for AWS Organizations in the *AWS Organizations User Guide*.

# AWS Outposts endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | outposts.us-east-2.amazonaws.com | HTTPS |
| | | outposts-fips.us-east-2.api.aws | HTTPS |
| | | outposts-fips.us-east-2.amazonaws.com | HTTPS |
| | | outposts.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | outposts.us-east-1.amazonaws.com | HTTPS |
| | | outposts-fips.us-east-1.api.aws | HTTPS |
| | | outposts-fips.us-east-1.amazonaws.com | HTTPS |
| | | outposts.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | outposts.us-west-1.amazonaws.com | HTTPS |
| | | outposts-fips.us-west-1.api.aws | HTTPS |
| | | outposts-fips.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | outposts.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | outposts.us-west-2.amazonaws.com | HTTPS |
| | | outposts-fips.us-west-2.api.aws | HTTPS |
| | | outposts-fips.us-west-2.amazonaws.com | HTTPS |
| | | outposts.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | outposts.af-south-1.amazonaws.com | HTTPS |
| | | outposts.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | outposts.ap-east-1.amazonaws.com | HTTPS |
| | | outposts.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | outposts.ap-southeast-3.amazonaws.com | HTTPS |
| | | outposts.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | outposts.ap-south-1.amazonaws.com | HTTPS |
| | | outposts.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | outposts.ap-northeast-3.amazonaws.com | HTTPS |
| | | outposts.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | outposts.ap-northeast-2.amazonaws.com | HTTPS |
| | | outposts.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | outposts.ap-southeast-1.amazonaws.com | HTTPS |
| | | outposts.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | outposts.ap-southeast-2.amazonaws.com | HTTPS |
| | | outposts.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | outposts.ap-northeast-1.amazonaws.com | HTTPS |
| | | outposts.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | outposts.ca-central-1.amazonaws.com | HTTPS |
| | | outposts-fips.ca-central-1.api.aws | HTTPS |
| | | outposts-fips.ca-central-1.amazonaws.com | HTTPS |
| | | outposts.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | outposts.eu-central-1.amazonaws.com | HTTPS |
| | | outposts.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | outposts.eu-west-1.amazonaws.com | HTTPS |
| | | outposts.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | outposts.eu-west-2.amazonaws.com | HTTPS |
| | | outposts.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | outposts.eu-south-1.amazonaws.com | HTTPS |
| | | outposts.eu-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | outposts.eu-west-3.amazonaws.com | HTTPS |
| | | outposts.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | outposts.eu-south-2.amazonaws.com | HTTPS |
| | | outposts.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | outposts.eu-north-1.amazonaws.com | HTTPS |
| | | outposts.eu-north-1.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | outposts.il-central-1.amazonaws.com | HTTPS |
| | | outposts.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | outposts.mx-central-1.amazonaws.com | HTTPS |
| | | outposts.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | outposts.me-south-1.amazonaws.com | HTTPS |
| | | outposts.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | outposts.me-central-1.amazonaws.com | HTTPS |
| | | outposts.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | outposts.sa-east-1.amazonaws.com | HTTPS |
| | | outposts.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | outposts.us-gov-east-1.amazonaws.com | HTTPS |
| | | outposts.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | outposts.us-gov-west-1.amazonaws.com<br>outposts.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

## Amazon S3 on Outposts

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | s3-outposts.us-east-2.amazonaws.com<br>s3-outposts-fips.us-east-2.api.aws<br>s3-outposts-fips.us-east-2.amazonaws.com<br>s3-outposts.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | s3-outposts.us-east-1.amazonaws.com<br>s3-outposts-fips.us-east-1.api.aws<br>s3-outposts-fips.us-east-1.amazonaws.com<br>s3-outposts.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | s3-outposts.us-west-1.amazonaws.com<br>s3-outposts-fips.us-west-1.api.aws<br>s3-outposts-fips.us-west-1.amazonaws.com<br>s3-outposts.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | s3-outposts.us-west-2.amazonaws.com<br>s3-outposts-fips.us-west-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | s3-outposts-fips.us-west-2.amazonaws.com | HTTPS |
| | | s3-outposts.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | s3-outposts.af-south-1.amazonaws.com | HTTPS |
| | | s3-outposts.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | s3-outposts.ap-east-1.amazonaws.com | HTTPS |
| | | s3-outposts.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | s3-outposts.ap-southeast-3.amazonaws.com | HTTPS |
| | | s3-outposts.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | s3-outposts.ap-south-1.amazonaws.com | HTTPS |
| | | s3-outposts.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | s3-outposts.ap-northeast-3.amazonaws.com | HTTPS |
| | | s3-outposts.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | s3-outposts.ap-northeast-2.amazonaws.com | HTTPS |
| | | s3-outposts.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | s3-outposts.ap-southeast-1.amazonaws.com | HTTPS |
| | | s3-outposts.ap-southeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | s3-outposts.ap-southeast-2.amazonaws.com | HTTPS |
| | | s3-outposts.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | s3-outposts.ap-northeast-1.amazonaws.com | HTTPS |
| | | s3-outposts.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | s3-outposts.ca-central-1.amazonaws.com | HTTPS |
| | | s3-outposts-fips.ca-central-1.api.aws | HTTPS |
| | | s3-outposts-fips.ca-central-1.amazonaws.com | HTTPS |
| | | s3-outposts.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | s3-outposts.eu-central-1.amazonaws.com | HTTPS |
| | | s3-outposts.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | s3-outposts.eu-west-1.amazonaws.com | HTTPS |
| | | s3-outposts.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | s3-outposts.eu-west-2.amazonaws.com | HTTPS |
| | | s3-outposts.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | s3-outposts.eu-south-1.amazonaws.com | HTTPS |
| | | s3-outposts.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | s3-outposts.eu-west-3.amazonaws.com | HTTPS |
| | | s3-outposts.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | s3-outposts.eu-north-1.amazonaws.com<br>s3-outposts.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | s3-outposts.il-central-1.amazonaws.com<br>s3-outposts.il-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | s3-outposts.me-south-1.amazonaws.com<br>s3-outposts.me-south-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | s3-outposts.sa-east-1.amazonaws.com<br>s3-outposts.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | s3-outposts.us-gov-east-1.amazonaws.com<br>s3-outposts-fips.us-gov-east-1.api.aws<br>s3-outposts-fips.us-gov-east-1.amazonaws.com<br>s3-outposts.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | s3-outposts.us-gov-west-1.amazonaws.com<br>s3-outposts-fips.us-gov-west-1.api.aws<br>s3-outposts-fips.us-gov-west-1.amazonaws.com<br>s3-outposts.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Outpost sites | Each supported Region: 100 | Yes | The maximum number of Outpost sites that you can create in this account in the current region. An Outpost site is the customer-managed physical building where you power and attach your Outpost equipment to the network. |
| Outposts per site | Each supported Region: 10 | Yes | The maximum number of Outposts that you can create per site. AWS Outposts includes hardware and virtual resources known as Outposts. This quota limits your Outpost virtual resources. |

# AWS Panorama endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Applications per device | Each supported Region: 10 | No | The maximum number of application instances that you can run on a device. |
| Assets per package | Each supported Region: 50 | No | The maximum number of assets that you upload per node package. |
| Camera nodes per application | Each supported Region: 8 | No | The maximum number of camera nodes that you can add to an applicati on. |
| Devices | Each supported Region: 50 | No | The maximum number of devices that you can register in this account in the current AWS Region. |
| Node packages | Each supported Region: 50 | No | The maximum number of node packages that you can create in this account in the current AWS Region. |
| Package artifact size | Each supported Region: 8 Gigabytes | No | The maximum size of a node package artifact. |
| Rate of CreateApplicationInstance API requests | Each supported Region: 1 | No | The maximum number of CreateApplicationI |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | nstance API requests per second that you can perform in this account in the current AWS Region. |
| Rate of CreateJobForDevices API requests | Each supported Region: 1 | No | The maximum number of CreateJobForDevices API requests per second that you can perform in this account in the current AWS Region. |
| Rate of CreateNodeFromTemplateJob API requests | Each supported Region: 1 | No | The maximum number of CreateNodeFromTemplateJob API requests per second that you can perform in this account in the current AWS Region. |
| Rate of CreatePackage API requests | Each supported Region: 1 | No | The maximum number of CreatePackage API requests per second that you can perform in this account in the current AWS Region. |
| Rate of CreatePackageImportJob API requests | Each supported Region: 1 | No | The maximum number of CreatePackageImportJob API requests per second that you can perform in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteDevice API requests | Each supported Region: 1 | No | The maximum number of DeleteDevice API requests per second that you can perform in this account in the current AWS Region. |
| Rate of DeletePackage API requests | Each supported Region: 1 | No | The maximum number of DeletePackage API requests per second that you can perform in this account in the current AWS Region. |
| Rate of DeregisterPackageVersion API requests | Each supported Region: 1 | No | The maximum number of DeregisterPackageVersion API requests per second that you can perform in this account in the current AWS Region. |
| Rate of DescribePackage API requests | Each supported Region: 1 | No | The maximum number of DescribePackage API requests per second that you can perform in this account in the current AWS Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ProvisionDevice API requests | Each supported Region: 1 | No | The maximum number of ProvisionDevice API requests per second that you can perform in this account in the current AWS Region. |
| Rate of RegisterPackageVersion API requests | Each supported Region: 1 | No | The maximum number of RegisterPackageVersion API requests per second that you can perform in this account in the current AWS Region. |
| Rate of RemoveApplicationInstance API requests | Each supported Region: 1 | No | The maximum number of RemoveApplicationInstance API requests per second that you can perform in this account in the current AWS Region. |
| Rate of deployments | Each supported Region: 200 | No | The maximum number of times that you can deploy an application per day, per device. |
| Rate of total API requests | Each supported Region: 5 | No | The maximum number of API requests per second that you can perform in this account in the current AWS Region. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Versions per package | Each supported Region: 20 | No | The maximum number of package versions that you can create per node package. |

# AWS Payment Cryptography endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Control plane endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | controlplane.payment-cryptography.us-east-2.amazonaws.com<br><br>controlplane.payment-cryptography.us-east-2.api.aws | HTTPS<br><br>HTTPS | |
| US East (N. Virginia) | us-east-1 | controlplane.payment-cryptography.us-east-1.amazonaws.com<br><br>controlplane.payment-cryptography.us-east-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | controlplane.payment-cryptography.us-west-2.amazonaws.com<br><br>controlplane.payment-cryptography.us-west-2.api.aws | HTTPS<br><br>HTTPS | |
| Africa (Cape Town) | af-south-1 | controlplane.payment-cryptography.af-south-1.amazonaws.com<br><br>controlplane.payment-cryptography.af-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | controlplane.payment-cryptography.ap-south-2.amazonaws.com<br><br>controlplane.payment-cryptography.ap-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | controlplane.payment-cryptography.ap-south-1.amazonaws.com<br><br>controlplane.payment-cryptography.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | controlplane.payment-cryptography.ap-northeast-3.amazonaws.com<br><br>controlplane.payment-cryptography.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | controlplane.payment-cryptography.ap-southeast-1.amazonaws.com<br><br>controlplane.payment-cryptography.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | controlplane.payment-cryptography.ap-southeast-2.amazonaws.com<br><br>controlplane.payment-cryptography.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | controlplane.payment-cryptography.ap-northeast-1.amazonaws.com<br><br>controlplane.payment-cryptography.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | controlplane.payment-cryptography.ca-central-1.amazonaws.com<br><br>controlplane.payment-cryptography.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | controlplane.payment-cryptography.eu-central-1.amazonaws.com<br><br>controlplane.payment-cryptography.eu-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | controlplane.payment-cryptography.eu-west-1.amazonaws.com<br><br>controlplane.payment-cryptography.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (London) | eu-west-2 | controlplane.payment-cryptography.eu-west-2.amazonaws.com<br><br>controlplane.payment-cryptography.eu-west-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | controlplane.payment-cryptography.eu-west-3.amazonaws.com<br><br>controlplane.payment-cryptography.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | controlplane.payment-cryptography.sa-east-1.amazonaws.com<br><br>controlplane.payment-cryptography.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## Data plane endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | dataplane.payment-cryptography.us-east-2.amazonaws.com<br><br>dataplane.payment-cryptography.us-east-2.api.aws | HTTPS<br><br>HTTPS | |
| US East (N. Virginia) | us-east-1 | dataplane.payment-cryptography.us-east-1.amazonaws.com<br><br>dataplane.payment-cryptography.us-east-1.api.aws | HTTPS<br><br>HTTPS | |
| US West (Oregon) | us-west-2 | dataplane.payment-cryptography.us-west-2.amazonaws.com<br><br>dataplane.payment-cryptography.us-west-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | dataplane.payment-cryptography.af-south-1.amazonaws.com<br><br>dataplane.payment-cryptography.af-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | dataplane.payment-cryptography.ap-south-2.amazonaws.com<br><br>dataplane.payment-cryptography.ap-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | dataplane.payment-cryptography.ap-south-1.amazonaws.com<br><br>dataplane.payment-cryptography.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | dataplane.payment-cryptography.ap-northeast-3.amazonaws.com<br><br>dataplane.payment-cryptography.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | dataplane.payment-cryptography.ap-southeast-1.amazonaws.com<br><br>dataplane.payment-cryptography.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | dataplane.payment-cryptography.ap-southeast-2.amazonaws.com<br><br>dataplane.payment-cryptography.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | dataplane.payment-cryptography.ap-northeast-1.amazonaws.com<br><br>dataplane.payment-cryptography.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | dataplane.payment-cryptography.ca-central-1.amazonaws.com<br><br>dataplane.payment-cryptography.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | dataplane.payment-cryptography.eu-central-1.amazonaws.com<br><br>dataplane.payment-cryptography.eu-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | dataplane.payment-cryptography.eu-west-1.amazonaws.com<br><br>dataplane.payment-cryptography.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (London) | eu-west-2 | dataplane.payment-cryptography.eu-west-2.amazonaws.com<br><br>dataplane.payment-cryptography.eu-west-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Paris) | eu-west-3 | dataplane.payment-cryptography.eu-west-3.amazonaws.com<br><br>dataplane.payment-cryptography.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | dataplane.payment-cryptography.sa-east-1.amazonaws.com<br><br>dataplane.payment-cryptography.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Aliases | Each supported Region: 2,000 | Yes | The maximum number of aliases you can have in this account in the current Region. |
| Combined rate of control plane requests | Each supported Region: 5 per second | Yes | The maximum number of control plane requests per second that you can make in this account in the current Region. This quota applies to all control plane operations combined. |
| Combined rate of data plane requests (asymmetric) | Each supported Region: 20 per second | Yes | The maximum number of requests per second for data plane operations with an asymmetric key that you can make in this account in the current Region. This quota applies to all |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | data plane operations combined. |
| Combined rate of data plane requests (symmetric) | Each supported Region: 500 per second | Yes | The maximum number of requests per second for data plane operations with a symmetric key that you can make in this account in the current Region. This quota applies to all data plane operations combined. |
| Keys | Each supported Region: 2,000 | Yes | The maximum number of keys you can have in this account in the current Region, excluding deleted keys. |

# AWS Parallel Computing Service endpoints and quotas

The following tables describe the service endpoints and service quotas for AWS Parallel Computing Service. Service quotas, also referred to as *limits*, are the maximum number of service resources or operations for your AWS account. For more information, see *AWS service quotas*.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | pcs.us-east-2.amazonaws.com<br>pcs-fips.us-east-2.amazonaws.com<br>pcs-fips.us-east-2.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | pcs.us-east-2.api.aws | | |
| US East (N. Virginia) | us-east-1 | pcs.us-east-1.amazonaws.com<br>pcs-fips.us-east-1.amazonaws.com<br>pcs-fips.us-east-1.api.aws<br>pcs.us-east-1.api.aws | | |
| US West (Oregon) | us-west-2 | pcs.us-west-2.amazonaws.com<br>pcs-fips.us-west-2.amazonaws.com<br>pcs-fips.us-west-2.api.aws<br>pcs.us-west-2.api.aws | | |
| Africa (Cape Town) | af-south-1 | pcs.af-south-1.amazonaws.com<br>pcs.af-south-1.api.aws | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | pcs.ap-southeast-3.amazonaws.com<br>pcs.ap-southeast-3.api.aws | | |
| Asia Pacific (Mumbai) | ap-south-1 | pcs.ap-south-1.amazonaws.com<br>pcs.ap-south-1.api.aws | | |
| Asia Pacific (New Zealand) | ap-southeast-6 | pcs.ap-southeast-6.amazonaws.com<br>pcs.ap-southeast-6.api.aws | | |
| Asia Pacific (Osaka) | ap-northeast-3 | pcs.ap-northeast-3.amazonaws.com<br>pcs.ap-northeast-3.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | pcs.ap-northeast-2.amazonaws.com<br>pcs.ap-northeast-2.api.aws | | |
| Asia Pacific (Singapore) | ap-southeast-1 | pcs.ap-southeast-1.amazonaws.com<br>pcs.ap-southeast-1.api.aws | | |
| Asia Pacific (Sydney) | ap-southeast-2 | pcs.ap-southeast-2.amazonaws.com<br>pcs.ap-southeast-2.api.aws | | |
| Asia Pacific (Tokyo) | ap-northeast-1 | pcs.ap-northeast-1.amazonaws.com<br>pcs.ap-northeast-1.api.aws | | |
| Europe (Frankfurt) | eu-central-1 | pcs.eu-central-1.amazonaws.com<br>pcs.eu-central-1.api.aws | | |
| Europe (Ireland) | eu-west-1 | pcs.eu-west-1.amazonaws.com<br>pcs.eu-west-1.api.aws | | |
| Europe (London) | eu-west-2 | pcs.eu-west-2.amazonaws.com<br>pcs.eu-west-2.api.aws | | |
| Europe (Milan) | eu-south-1 | pcs.eu-south-1.amazonaws.com<br>pcs.eu-south-1.api.aws | | |
| Europe (Paris) | eu-west-3 | pcs.eu-west-3.amazonaws.com<br>pcs.eu-west-3.api.aws | | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | pcs.eu-south-2.amazonaws.com<br>pcs.eu-south-2.api.aws | | |
| Europe (Stockholm) | eu-north-1 | pcs.eu-north-1.amazonaws.com<br>pcs.eu-north-1.api.aws | | |
| South America (São Paulo) | sa-east-1 | pcs.sa-east-1.amazonaws.com<br>pcs.sa-east-1.api.aws | | |
| AWS GovCloud (US-East) | us-gov-east-1 | pcs.us-gov-east-1.amazonaws.com<br>pcs-fips.us-gov-east-1.amazonaws.com<br>pcs-fips.us-gov-east-1.api.aws<br>pcs.us-gov-east-1.api.aws | | |
| AWS GovCloud (US-West) | us-gov-west-1 | pcs.us-gov-west-1.amazonaws.com<br>pcs-fips.us-gov-west-1.amazonaws.com<br>pcs-fips.us-gov-west-1.api.aws<br>pcs.us-gov-west-1.api.aws | | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Clusters | Each supported Region: 5 | Yes | The maximum number of clusters per Region. |

## Internal quotas

The following quotas are internal and non-adjustable.

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrent cluster creation | 1 | No | The maximum number of clusters in the `Creating` state per AWS Region. |

## Relevant quotas for other AWS services

AWS PCS uses other AWS services. Your service quotas for those services impact your use of AWS PCS.

**Amazon EC2 service quotas that impact AWS PCS**

- Spot instance requests
- Running on-demand instances
- Launch templates
- Launch template versions
- Amazon EC2 API requests

For more information, see Amazon EC2 endpoints and quotas.

# Amazon Personalize endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Amazon Personalize

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | personalize.us-east-1.amazonaws.com<br><br>personalize.us-east-1.api.aws | HTTPS |
| US East (Ohio) | us-east-2 | personalize.us-east-2.amazonaws.com<br><br>personalize.us-east-2.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | personalize.us-west-2.amazonaws.com<br><br>personalize.us-west-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | personalize.ap-northeast-1.amazonaws.com<br><br>personalize.ap-northeast-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | personalize.ap-south-1.amazonaws.com<br><br>personalize.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | personalize.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | personalize.ap-northeast-2.api.aws | |
| Asia Pacific (Singapore) | ap-southeast-1 | personalize.ap-southeast-1.amazonaws.com<br><br>personalize.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | personalize.ap-southeast-2.amazonaws.com<br><br>personalize.ap-southeast-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | personalize.ca-central-1.amazonaws.com<br><br>personalize.ca-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | personalize.eu-west-1.amazonaws.com<br><br>personalize.eu-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | personalize.eu-central-1.amazonaws.com<br><br>personalize.eu-central-1.api.aws | HTTPS |

## Amazon Personalize Events

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | personalize-events.us-east-1.amazonaws.com<br><br>personalize-events.us-east-1.api.aws | HTTPS |
| US East (Ohio) | us-east-2 | personalize-events.us-east-2.amazonaws.com<br><br>personalize-events.us-east-2.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | personalize-events.us-west-2.amazonaws.com<br><br>personalize-events.us-west-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | personalize-events.ap-northeast-1.amazonaws.com<br><br>personalize-events.ap-northeast-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | personalize-events.ap-south-1.amazonaws.com<br><br>personalize-events.ap-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | personalize-events.ap-northeast-2.amazonaws.com<br><br>personalize-events.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | personalize-events.ap-southeast-1.amazonaws.com<br><br>personalize-events.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | personalize-events.ap-southeast-2.amazonaws.com<br><br>personalize-events.ap-southeast-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | personalize-events.ca-central-1.amazonaws.com<br><br>personalize-events.ca-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | personalize-events.eu-west-1.amazonaws.com<br><br>personalize-events.eu-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Frankfurt) | eu-central-1 | personalize-events.eu-central-1.amazonaws.com<br><br>personalize-events.eu-central-1.api.aws | HTTPS |

## Amazon Personalize Runtime

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | personalize-runtime.us-east-1.amazonaws.com<br><br>personalize-runtime.us-east-1.api.aws | HTTPS |
| US East (Ohio) | us-east-2 | personalize-runtime.us-east-2.amazonaws.com<br><br>personalize-runtime.us-east-2.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | personalize-runtime.us-west-2.amazonaws.com<br><br>personalize-runtime.us-west-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | personalize-runtime.ap-northeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | personalize-runtime.ap-northeast-1.api.aws | |
| Asia Pacific (Mumbai) | ap-south-1 | personalize-runtime.ap-south-1.amazonaws.com<br><br>personalize-runtime.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | personalize-runtime.ap-northeast-2.amazonaws.com<br><br>personalize-runtime.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | personalize-runtime.ap-southeast-1.amazonaws.com<br><br>personalize-runtime.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | personalize-runtime.ap-southeast-2.amazonaws.com<br><br>personalize-runtime.ap-southeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | personalize-runtime.ca-central-1.amazonaws.com<br><br>personalize-runtime.ca-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | personalize-runtime.eu-west-1.amazonaws.com<br><br>personalize-runtime.eu-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | personalize-runtime.eu-central-1.amazonaws.com<br><br>personalize-runtime.eu-central-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Active campaigns per dataset group | Each supported Region: 5 | Yes | The total number of active campaigns per dataset group in the current Region. |
| Active dataset groups | Each supported Region: 5 | Yes | The total number of active dataset groups that you can create |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | in this account in the current Region. |
| Active filters per dataset group | Each supported Region: 30 | Yes | The total number of active filters per dataset group in the current Region. |
| Active solutions per dataset group | Each supported Region: 20 | Yes | The total number of active solutions per dataset group in the current Region. |
| Amount of data for HRNN recipe | Each supported Region: 100 Gigabytes | No | The maximum amount of data for an individual dataset for HRNN recipe |
| Amount of data for Personalized-Ranking recipe | Each supported Region: 100 Gigabytes | No | The maximum amount of data for an individual dataset for Personalized-Ranking recipe |
| Amount of data for Popularity-Count recipe | Each supported Region: 100 Gigabytes | No | The maximum amount of data for an individual dataset for Popularity-Count recipe |
| Amount of data for SIMS recipe | Each supported Region: 100 Gigabytes | No | The maximum amount of data for an individual dataset for SIMS recipe |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amount of data per incremental import. | Each supported Region: 1 Gigabytes | Yes | The maximum amount of data (in GB) you can import with a single incremental import in this account in the current Region. |
| Amount of interactions data for HRNN-coldstart recipe | Each supported Region: 100 Gigabytes | No | The maximum amount of data for interactions dataset for HRNN-cold start recipe |
| Amount of interactions data for HRNN-metadata recipe | Each supported Region: 100 Gigabytes | No | The maximum amount of data for interactions dataset for HRNN-meta data recipe |
| Amount of users and items data combined for HRNN-coldstart recipe | Each supported Region: 5 Gigabytes | No | The maximum amount of data for users dataset and items dataset combined for HRNN-cold start recipe |
| Amount of users and items data combined for HRNN-metadata recipe | Each supported Region: 5 Gigabytes | No | The maximum amount of data for users dataset and items dataset combined for HRNN-metadata recipe |
| Event size | Each supported Region: 10 Kilobytes | No | The maximum size of an event in kilobytes. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of action interactions per event type per user considered by a filter. | Each supported Region: 300 | Yes | The maximum number of action interactions per event type per user Amazon Personalize considers when filtering recommendations. This quota applies to each get action recommendations request in this account in the current region. |
| Maximum number of interactions per event type per user considered by a filter. | Each supported Region: 100 | Yes | The maximum number of interactions per event type per user Amazon Personalize considers when filtering recommendations. This quota applies to each get recommendations request in this account in the current Region. |
| Maximum number of recommenders per dataset group | Each supported Region: 5 | No | The total number of recommenders per dataset group in the current Region. |
| Minimum data points for model training | Each supported Region: 1,000 | No | The minimum number of data points required for training a model (creating a solution) |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum unique users for model training | Each supported Region: 25 | No | The minimum number of unique users required for training a model (creating a solution). |
| Number of action interaction events in a PutActionInteractions call | Each supported Region: 10 | No | Maximum number of action interaction events in a PutActionInteractions call. |
| Number of action interactions for model training | Each supported Region: 500,000,000 | No | Maximum number of action interactions that are considered by a model during training. |
| Number of actions used in model training | Each supported Region: 1,000 | No | Maximum number of actions that are considered by a model during training and generating recommendations. |
| Number of events in PutEvents call | Each supported Region: 10 | No | The maximum number of events in a PutEvents call. |
| Number of interactions for model training | Each supported Region: 500,000,000 | Yes | The maximum number of interactions that are considered by a model during training. |
| Number of items used in model training | Each supported Region: 750,000 | No | The maximum number of items that are considered by a model during training. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of schemas | Each supported Region: 500 | No | The total number of active schemas that you can create in this account in the current Region. |
| Pending data deletion jobs per dataset group | Each supported Region: 5 | Yes | The total number of pending data deletion jobs that you can create per dataset group in the current Region. |
| Pending or In Progress batch inference jobs | Each supported Region: 5 | Yes | The total number of pending or in progress batch inference jobs that you can create in this account in the current Region. |
| Pending or In Progress solution versions | Each supported Region: 20 | Yes | The total number of pending or in progress solution versions that you can create in this account in the current Region. |
| PutActionInteractions Event size | Each supported Region: 10 Kilobytes | No | The maximum size of a PutActionInteractions event in kilobytes. |
| Rate of CreateCampaign requests | Each supported Region: 1 | No | The maximum number of CreateCampaign requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateDataset requests | Each supported Region: 1 | No | The maximum number of CreateDataset requests that you can make per second. |
| Rate of CreateDatasetGroup requests | Each supported Region: 1 | No | The maximum number of CreateDatasetGroup requests that you can make per second |
| Rate of CreateDatasetImportJob requests | Each supported Region: 1 | No | The maximum number of CreateDatasetImportJob requests that you can make per second. |
| Rate of CreateEventTracker requests | Each supported Region: 1 | No | The maximum number of CreateEventTracker requests that you can make per second. |
| Rate of CreateSchema requests | Each supported Region: 1 | No | The maximum number of CreateSchema requests that you can make per second. |
| Rate of CreateSolution requests | Each supported Region: 1 | No | The maximum number of CreateSolution requests that you can make per second. |
| Rate of CreateSolutionVersion requests | Each supported Region: 1 | No | The maximum number of CreateSolutionVersion requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DeleteCampaign requests | Each supported Region: 1 | No | The maximum number of DeleteCampaign requests that you can make per second. |
| Rate of DeleteDataset requests | Each supported Region: 1 | No | The maximum number of DeleteDataset requests that you can make per second. |
| Rate of DeleteDatasetGroup requests | Each supported Region: 1 | No | The maximum number of DeleteDatasetGroup requests that you can make per second. |
| Rate of DeleteDatasetImportJob requests | Each supported Region: 1 | No | The maximum number of DeleteDatasetImportJob requests that you can make per second. |
| Rate of DeleteEventTracker requests | Each supported Region: 1 | No | The maximum number of DeleteEventTracker requests that you can make per second. |
| Rate of DeleteSchema requests | Each supported Region: 1 | No | The maximum number of DeleteSchema requests that you can make per second |
| Rate of DeleteSolution requests | Each supported Region: 1 | No | The maximum number of DeleteSolution requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DescribeAlgorithm requests | Each supported Region: 1 | No | The maximum number of DescribeAlgorithm requests that you can make per second. |
| Rate of DescribeCampaign requests | Each supported Region: 1 | No | The maximum number of DescribeCampaign requests that you can make per second. |
| Rate of DescribeDataset requests | Each supported Region: 1 | No | The maximum number of DescribeDataset requests that you can make per second. |
| Rate of DescribeDatasetGroup requests | Each supported Region: 1 | No | The maximum number of DescribeDatasetGroup requests that you can make per second. |
| Rate of DescribeDatasetImportJob requests | Each supported Region: 1 | No | The maximum number of DescribeDatasetImportJob requests that you can make per second. |
| Rate of DescribeEventTracker requests | Each supported Region: 1 | No | The maximum number of DescribeEventTracker requests that you can make per second. |
| Rate of DescribeFeatureTransformation requests | Each supported Region: 1 | No | The maximum number of DescribeFeatureTransformation requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DescribeRecipe requests | Each supported Region: 1 | No | The maximum number of DescribeRecipe requests that you can make per second. |
| Rate of DescribeSchema requests | Each supported Region: 1 | No | The maximum number of DescribeSchema requests that you can make per second. |
| Rate of DescribeSolution requests | Each supported Region: 1 | No | The maximum number of DescribeSolution requests that you can make per second. |
| Rate of GetActionRecommendations requests per campaign | Each supported Region: 500 | No | The maximum number of GetActionRecommendations requests that you can make per second per campaign. |
| Rate of GetPersonalizedRanking requests per campaign | Each supported Region: 500 | No | The maximum number of GetPersonalizedRanking requests that you can make per second per campaign. |
| Rate of GetRecommendations requests per campaign | Each supported Region: 500 | No | The maximum number of GetRecommendations requests that you can make per second per campaign. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetSolutionMetrics requests | Each supported Region: 1 | No | The maximum number of GetSolutionMetrics requests that you can make per second. |
| Rate of ListCampaigns requests | Each supported Region: 1 | No | The maximum number of ListCampaigns requests that you can make per second. |
| Rate of ListDatasetGroups requests | Each supported Region: 1 | No | The maximum number of ListDatasetGroups requests that you can make per second. |
| Rate of ListDatasetImportJobRuns requests | Each supported Region: 1 | No | The maximum number of ListDatasetImportJobRuns requests that you can make per second. |
| Rate of ListDatasetImportJobs requests | Each supported Region: 1 | No | The maximum number of ListDatasetImportJobs requests that you can make per second. |
| Rate of ListDatasets requests | Each supported Region: 1 | No | The maximum number of ListDatasets requests that you can make per second. |
| Rate of ListEventTrackers requests | Each supported Region: 1 | No | The maximum number of ListEventTrackers requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListRecipes requests | Each supported Region: 1 | No | The maximum number of ListRecipes requests that you can make per second. |
| Rate of ListSchemas requests | Each supported Region: 1 | No | The maximum number of ListSchemas requests that you can make per second. |
| Rate of ListSolutionVersions requests | Each supported Region: 1 | No | The maximum number of ListSolutionVersions requests that you can make per second. |
| Rate of ListSolutions requests | Each supported Region: 1 | No | The maximum number of ListSolutions requests that you can make per second. |
| Rate of PutActionInteractions requests per dataset group | Each supported Region: 1,000 | Yes | The maximum number of PutActionInteractions requests that you can make per second per dataset group from this account in the current region. |
| Rate of PutEvents requests per dataset group | Each supported Region: 1,000 | Yes | The maximum number of PutEvents requests that you can make per second per dataset group from this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateCampaign requests | Each supported Region: 1 | No | The maximum number of UpdateCampaign requests that you can make per second. |
| Rate of UpdateDataset requests | Each supported Region: 1 | No | The maximum number of UpdateDataset requests that you can make per second. |
| Rate of transactions per account | Each supported Region: 2,500 | No | The maximum transactions per second per account. Transaction is a single recommendation request. |

## Amazon Pinpoint endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

Amazon Pinpoint includes the Amazon Pinpoint API and the Amazon Pinpoint SMS and Voice API.

# Service endpoints

## Amazon Pinpoint API

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | pinpoint.us-east-2.amazonaws.com<br>pinpoint.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | pinpoint.us-east-1.amazonaws.com<br>pinpoint.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | pinpoint.us-west-2.amazonaws.com<br>pinpoint.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | pinpoint.ap-south-1.amazonaws.com<br>pinpoint.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | pinpoint.ap-northeast-2.amazonaws.com<br>pinpoint.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | pinpoint.ap-southeast-1.amazonaws.com<br>pinpoint.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | pinpoint.ap-southeast-2.amazonaws.com<br>pinpoint.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | pinpoint.ap-northeast-1.amazonaws.com<br>pinpoint.ap-northeast-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | pinpoint.ca-central-1.amazonaws.com<br>pinpoint.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | pinpoint.eu-central-1.amazonaws.com<br>pinpoint.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | pinpoint.eu-west-1.amazonaws.com<br>pinpoint.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | pinpoint.eu-west-2.amazonaws.com<br>pinpoint.eu-west-2.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | pinpoint.us-gov-west-1.amazonaws.com<br>pinpoint.us-gov-west-1.api.aws | HTTPS | |

## Amazon Pinpoint SMS and Voice API v1

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | sms-voice.pinpoint.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | sms-voice.pinpoint.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | sms-voice.pinpoint.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | sms-voice.pinpoint.ap-southeast-2.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | sms-voice.pinpoint.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | sms-voice.pinpoint.eu-west-1.amazonaws.com | HTTPS | |

**Note**

The Amazon Pinpoint SMS and Voice API is not available in the following Regions:

- Asia Pacific (Seoul) Region
- Asia Pacific (Singapore) Region
- Asia Pacific (Tokyo) Region
- Canada (Central) Region
- Europe (London) Region

## Amazon Pinpoint SMS and Voice API v2

For AWS End User Messaging SMS and Voice V2 API service endpoints see AWS End User Messaging endpoints and quotas.

# Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| APNs sandbox message payload size per message | Each supported Region: 4 Kilobytes | No | The maximum APNs sandbox message payload size (in KB) per message. |
| Active campaigns per account | Each supported Region: 200 | No | The maximum number of active campaigns per account. An active campaign is a campaign that hasnt completed or failed. Active campaigns have a status of SCHEDULED, EXECUTING, or PENDING_NEXT_RUN. |
| Active in-app campaigns per project | Each supported Region: 25 | No | The maximum number of active in-app campaigns per project. |
| All other operations burst quota | Each supported Region: 300 | No | The maximum number of all other operation requests that you can make at one time. |
| All other operations rate quota | Each supported Region: 300 | No | The maximum number of all other operation requests that you can make per second. |
| Amazon Device Messaging (ADM) message payload size per message | Each supported Region: 6 Kilobytes | No | The maximum Amazon Device Messaging (ADM) |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | message payload size (in KB) per message. |
| Apple Push Notification service (APNs) message payload size per message | Each supported Region: 4 Kilobytes | No | The maximum Apple Push Notification service (APNs) message payload size (in KB) per message. |
| Attribute name length | Each supported Region: 50 | No | The maximum attribute name length in characters. |
| Attribute value length | Each supported Region: 100 | No | The maximum attribute value length in characters. |
| Baidu Cloud Push message payload size per message | Each supported Region: 4 Kilobytes | No | The maximum Baidu Cloud Push message payload size (in KB) per message. |
| CreateCampaign operation burst quota | Each supported Region: 25 | No | The maximum number of CreateCampaign operation requests that you can make at one time. |
| CreateCampaign operation rate quota | Each supported Region: 25 | No | The maximum number of CreateCampaign operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateEmailTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of CreateEmailTemplate operation requests that you can make at one time. |
| CreateEmailTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of CreateEmailTemplate operation requests that you can make per second. |
| CreateInAppTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of CreateInAppTemplate operation requests that you can make at one time. |
| CreateInAppTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of CreateInAppTemplate operation requests that you can make per second. |
| CreatePushTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of CreatePushTemplate operation requests that you can make at one time. |
| CreatePushTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of CreatePushTemplate operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CreateSegment operation burst quota | Each supported Region: 25 | No | The maximum number of CreateSegment operation requests that you can make at one time. |
| CreateSegment operation rate quota | Each supported Region: 25 | No | The maximum number of CreateSegment operation requests that you can make per second. |
| CreateSmsTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of CreateSmsTemplate operation requests that you can make at one time. |
| CreateSmsTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of CreateSmsTemplate operation requests that you can make per second. |
| CreateVoiceTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of CreateVoiceTemplate operation requests that you can make at one time. |
| CreateVoiceTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of CreateVoiceTemplate operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DeleteCampaign operation burst quota | Each supported Region: 25 | No | The maximum number of DeleteCampaign operation requests that you can make at one time. |
| DeleteCampaign operation rate quota | Each supported Region: 25 | No | The maximum number of DeleteCampaign operation requests that you can make per second. |
| DeleteEndpoint operation burst quota | Each supported Region: 5 | No | The maximum number of DeleteEndpoint operation requests that you can make at one time. |
| DeleteEndpoint operation rate quota | Each supported Region: 5 | No | The maximum number of DeleteEndpoint operation requests that you can make per second. |
| DeleteSegment operation burst quota | Each supported Region: 25 | No | The maximum number of DeleteSegment operation requests that you can make at one time. |
| DeleteSegment operation rate quota | Each supported Region: 25 | No | The maximum number of DeleteSegment operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Firebase Cloud Messaging (FCM) message payload size per message | Each supported Region: 4 Kilobytes | No | The maximum Firebase Cloud Messaging (FCM) message payload size (in KB) per message. |
| GetEndpoint operation burst quota | Each supported Region: 10 | No | The maximum number of GetEndpoint operation requests that you can make at one time. |
| GetEndpoint operation rate quota | Each supported Region: 10 | No | The maximum number of GetEndpoint operation requests that you can make per second. |
| Import size per import job | Each supported Region: 1 | No | The maximum import size (in GB) per import job. |
| Invocation payload size | Each supported Region: 7 Megabytes | No | The maximum invocation payload size (in MB). |
| Maximum amount of time to wait for a Lambda function to process data | Each supported Region: 15 Seconds | No | The maximum amount of time (in seconds) to wait for a Lambda function to process data. |
| Maximum length of a recommended attribute display name | Each supported Region: 25 | No | The maximum length of a recommended attribute display name (the name that appears in the Attribute finder on the console). |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum length of a recommended attribute name | Each supported Region: 50 | No | The maximum length of a recommended attribute name. |
| Maximum length of a recommended attribute value thats retrieved from Amazon Personalize | Each supported Region: 100 | No | Maximum length (in characters) of a recommended attribute value thats retrieved from Amazon Personalize. |
| Maximum message size, including attachments | Each supported Region: 10 Megabytes | No | The maximum message size (in MB), including attachments. |
| Maximum number of active event triggered journeys per account | Each supported Region: 20 | No | The maximum number of active event triggered journeys per account. |
| Maximum number of active journeys per account | Each supported Region: 50 | No | The maximum number of active journeys per account. |
| Maximum number of attempts to invoke a Lambda function | Each supported Region: 3 Seconds | No | The maximum number of attempts to invoke a Lambda function. |
| Maximum number of attribute keys and metric keys for each event per request | Each supported Region: 40 | No | The maximum number of attribute keys and metric keys for each event per request. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of characters for an event attribute name in a custom channel response | Each supported Region: 128 | No | The maximum number of characters for an event attribute name in a custom channel response. |
| Maximum number of characters for an event attribute value in a custom channel response | Each supported Region: 128 | No | The maximum number of characters for an event attribute value in a custom channel response. |
| Maximum number of characters in ADM-specific template parts of a push notification template | Each supported Region: 4,000 | No | The maximum number of characters in ADM-specific template parts of a push notification template. |
| Maximum number of characters in APN-specific template parts of a push notification template | Each supported Region: 2,000 | No | The maximum number of characters in APN-specific template parts of a push notification template. |
| Maximum number of characters in Baidu-specific template parts of a push notification template | Each supported Region: 4,000 | No | The maximum number of characters in Baidu-specific template parts of a push notification template. |
| Maximum number of characters in FCM-specific template parts of a push notification template | Each supported Region: 4,000 | No | The maximum number of characters in FCM-specific template parts of a push notification template. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of characters in a voice template | Each supported Region: 10,000 | No | The maximum number of characters in a voice template. |
| Maximum number of characters in an SMS template | Each supported Region: 1,600 | No | The maximum number of characters in an SMS template. |
| Maximum number of characters in an email template | Each supported Region: 500,000 | No | The maximum number of characters in an email template. |
| Maximum number of characters in the default template parts of a push notification template | Each supported Region: 2,000 | No | The maximum number of characters in the default template parts of a push notification template. |
| Maximum number of characters per attribute key | Each supported Region: 50 | No | The maximum number of characters per attribute key. |
| Maximum number of characters per attribute value | Each supported Region: 200 | No | The maximum number of characters per attribute value. If the number of characters exceeds 200 the event is dropped. |
| Maximum number of custom attribute keys per app | Each supported Region: 500 | No | The maximum number of custom attribute keys per app. |
| Maximum number of custom attribute values per attribute key | Each supported Region: 100,000 | No | The maximum number of custom attribute values per attribute key. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Maximum number of custom event types per app | Each supported Region: 1,500 | No | The maximum number of custom event types per app. |
| Maximum number of custom metric keys per app | Each supported Region: 500 | No | The maximum number of custom metric keys per app. |
| Maximum number of dimensions that can be used to create a segment | Each supported Region: 100 | No | The maximum number of dimensions that can be used to create a segment. |
| Maximum number of event attributes per endpoint in a custom channel response | Each supported Region: 5 | No | The maximum number of event attributes per endpoint in a custom channel response. |
| Maximum number of events in a request | Each supported Region: 100 | No | The maximum number of events in a request. |
| Maximum number of journey activities per journey | Each supported Region: 40 | No | The maximum number of journey activities per journey. |
| Maximum number of message templates per account | Each supported Region: 10,000 | No | The maximum number of message templates per account. |
| Maximum number of model configurations per account | Each supported Region: 100 | No | The maximum number of model configurations per account. |
| Maximum number of model configurations per message template | Each supported Region: 1 | No | The maximum number of model configurations per message template. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Maximum number of push notifications that can be sent per second in a campaign | Each supported Region: 25,000 | No | The maximum number of push notifications that can be sent per second in a campaign. |
| Maximum number of recommendations per endpoint or user | Each supported Region: 5 | No | The maximum number of recommendations per endpoint or user. |
| Maximum number of recommended attributes per endpoint or user | Each supported Region: 1 | No | The maximum number of recommended attributes per endpoint or user if the attribute values arent processed by an AWS Lambda function. |
| Maximum number of recommended attributes per endpoint or user (AWS Lambda function) | Each supported Region: 10 | No | The maximum number of recommended attributes per endpoint or user if the attribute values are processed by an AWS Lambda function. |
| Maximum number of versions per template | Each supported Region: 5,000 | No | The maximum number of versions per template. |
| Maximum segment size per campaign | Each supported Region: 100,000,000 | No | The maximum segment size for imported segments per campaign. For dynamic segments: unlimited. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum segment size per journey | Each supported Region: 100,000,000 | No | The maximum segment size per journey. For imported segments: 100,000,000 per journey. For dynamic segments: unlimited. |
| Maximum size of a request | Each supported Region: 4 Megabytes | No | The maximum size (in MB) of a request. |
| Maximum size of an individual event | Each supported Region: 1,000 Kilobytes | No | The maximum size (in KB) of an individual event. |
| Maximum size of an invocation payload (request and response) for a Lambda function | Each supported Region: 6 Megabytes | No | The maximum size (in MB) of an invocation payload (request and response) for a Lambda function. |
| Maximum size per endpoint | Each supported Region: 15 Kilobytes | No | The maximum size (in KB) per endpoint. |
| Number of Amazon Pinpoint projects | Each supported Region: 100 | No | The maximum number of Amazon Pinpoint projects per account. |
| Number of Amazon SNS topics for two-way SMS per account | Each supported Region: 100,000 | Yes | The maximum number of Amazon SNS topics for two-way SMS per account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of EndpointBatchItem objects in an EndpointBatchRequest payload | Each supported Region: 100 | No | The maximum number of EndpointBatchItem objects in an EndpointBatchRequest payload. The payload size cant exceed 7 MB. |
| Number of SMS messages that can be sent each second (sending rate) | Each supported Region: 20 | Yes | The maximum number of SMS messages that can be sent each second (sending rate). |
| Number of SMS messages that can be sent to a single recipient each second | Each supported Region: 1 | No | The maximum number of SMS messages that can be sent to a single recipient each second. |
| Number of attributes assigned to the Attributes parameter | Each supported Region: 250 | No | The maximum number of attributes assigned to the Attributes parameter. |
| Number of attributes assigned to the Attributes, Metrics, and UserAttributes parameters collectively | Each supported Region: 250 | No | The maximum number of attributes assigned to the Attributes, Metrics, and UserAttributes parameters collectively per endpoint. |
| Number of attributes assigned to the Metrics parameter | Each supported Region: 250 | No | The maximum number of attributes assigned to the Metrics parameter. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of attributes assigned to the UserAttributes parameter | Each supported Region: 250 | No | The maximum number of attributes assigned to the UserAttributes parameter. |
| Number of characters in a voice message | Each supported Region: 6,000 | No | The maximum number of characters in a voice message. 3,000 billable characters, in words that are spoken 6,000 characters total, including billable characters and SSML tags. |
| Number of concurrent import jobs | Each supported Region: 10 | No | The maximum number of concurrent import jobs per account. |
| Number of emails that can be sent each second (sending rate) | Each supported Region: 1 | Yes | The maximum number of emails that can be sent each second (sending rate). If your account is in the sandbox, 1 email per second. If your account is out of the sandbox, the rate varies based on your specific use case. This rate is based on the number of recipients, as opposed to the number of unique messages sent. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of emails that can be sent per 24-hour period (sending quota) | Each supported Region: 200 | Yes | The maximum number of emails that can be sent per 24-hour period (sending quota). If your account is in the sandbox, 200 emails per 24-hour period. If your account is out of the sandbox, the quota varies based on your specific use case. This quota is based on the number of recipients, as opposed to the number of unique messages sent. |
| Number of endpoints with the same user ID | Each supported Region: 10 | No | The maximum number of endpoints with the same user ID. |
| Number of event-based campaigns | Each supported Region: 25 | No | The maximum number of event-based campaigns . Campaigns that use event-based triggers have to use dynamic segments. They cant use imported segments. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of identities that you can verify | Each supported Region: 10,000 | No | The maximum number of identities that you can verify per AWS region. Identities refers to email addresses or domains, or any combination of the two. Every email you send using Amazon Pinpoint must be sent from a verified identity. |
| Number of import files per import job | Each supported Region: 10,000 | No | The maximum number of import files per import job. |
| Number of recipients per message | Each supported Region: 50 | No | The maximum number of recipients per message. |
| Number of values assigned to the Attributes parameter attributes per attribute | Each supported Region: 50 | No | The maximum number of values assigned to the Attributes parameter attributes per attribute. |
| Number of values assigned to the UserAttributes parameter attributes per attribute | Each supported Region: 50 | No | The maximum number of values assigned to the UserAttributes parameter attributes per attribute. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of verified identities | Each supported Region: 10,000 | No | The maximimum number of verified identities. Identities refers to email addresses or domains, or any combination of the two. Every email you send using Amazon Pinpoint must be sent from a verified identity. |
| Number of voice configuration sets per AWS region | Each supported Region: 10,000 | No | The maximum number of configuration sets per AWS region. |
| Number of voice messages that can be sent during a 24-hour period | Each supported Region: 20 | No | The maximum number of voice messages that can be sent during a 24-hour period. If your account is in the sandbox: 20 messages. If your account is out of the sandbox: unlimited. |
| Number of voice messages that can be sent from a single originating phone number per second | Each supported Region: 1 | No | The maximum number of voice messages that can be sent from a single originating phone number per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of voice messages that can be sent per minute | Each supported Region: 5 | No | The maximum number of voice messages that can be sent per minute. If your account is in the sandbox: 5 calls per minute. If your account is out of the sandbox: 20 calls per minute. |
| Number of voice messages that can be sent to a single recipient during a 24-hour period | Each supported Region: 5 | No | The maximum number of voice messages that can be sent to a single recipient during a 24-hour period. |
| PhoneNumberValidate operation burst quota | Each supported Region: 20 | No | The maximum number of PhoneNumberValidate operation requests that you can make at one time. |
| PhoneNumberValidate operation rate quota | Each supported Region: 20 | No | The maximum number of PhoneNumberValidate operation requests that you can make per second. |
| PutEvents operation burst quota | Each supported Region: 15 | No | The maximum number of PutEvents operation requests that you can make at one time. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| PutEvents operation rate quota | Each supported Region: 15 | No | The maximum number of PutEvents operation requests that you can make per second. |
| SMS spending threshold | Each supported Region: 1 | Yes | The maximum SMS spending threshold per account (in USD). |
| SendMessages operation burst quota | Each supported Region: 4,000 | No | The maximum number of SendMessages operation requests that you can make at one time. |
| SendMessages operation rate quota | Each supported Region: 4,000 | No | The maximum number of SendMessages operation requests that you can make per second. |
| SendUsersMessages operation burst quota | Each supported Region: 6,000 | No | The maximum number of SendUsersMessages operation requests that you can make at one time. |
| SendUsersMessages operation rate quota | Each supported Region: 6,000 | No | The maximum number of SendUsersMessages operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UpdateCampaign operation burst quota | Each supported Region: 25 | No | The maximum number of UpdateCampaign operation requests that you can make at one time. |
| UpdateCampaign operation rate quota | Each supported Region: 25 | No | The maximum number of UpdateCampaign operation requests that you can make per second. |
| UpdateEmailTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of UpdateEmailTemplate operation requests that you can make at one time. |
| UpdateEmailTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of UpdateEmailTemplate operation requests that you can make per second. |
| UpdateEndpoint operation burst quota | Each supported Region: 10 | No | The maximum number of UpdateEndpoint operation requests that you can make at one time. |
| UpdateEndpoint operation rate quota | Each supported Region: 10 | No | The maximum number of UpdateEndpoint operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| UpdateEndpointsBatch operation burst quota | Each supported Region: 2 | No | The maximum number of UpdateEndpointsBatch operation requests that you can make at one time. |
| UpdateEndpointsBatch operation rate quota | Each supported Region: 2 | No | The maximum number of UpdateEndpointsBatch operation requests that you can make per second. |
| UpdateInAppTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of UpdateInAppTemplate operation requests that you can make at one time. |
| UpdateInAppTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of UpdateInAppTemplate operation requests that you can make per second. |
| UpdatePushTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of UpdatePushTemplate operation requests that you can make at one time. |
| UpdatePushTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of UpdatePushTemplate operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| UpdateSegment operation burst quota | Each supported Region: 25 | No | The maximum number of UpdateSegment operation requests that you can make at one time. |
| UpdateSegment operation rate quota | Each supported Region: 25 | No | The maximum number of UpdateSegment operation requests that you can make per second. |
| UpdateSmsTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of UpdateSmsTemplate operation requests that you can make at one time. |
| UpdateSmsTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of UpdateSmsTemplate operation requests that you can make per second. |
| UpdateVoiceTemplate operation burst quota | Each supported Region: 10 | No | The maximum number of UpdateVoiceTemplate operation requests that you can make at one time. |
| UpdateVoiceTemplate operation rate quota | Each supported Region: 10 | No | The maximum number of UpdateVoiceTemplate operation requests that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Voice message length | Each supported Region: 30 | No | The maximum voice message length. If your account is in the sandbox: 30 seconds. If your account is out of the sandbox: 5 minutes. |

For more information, see Amazon Pinpoint quotas in the *Amazon Pinpoint Developer Guide*.

# Amazon Polly endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | polly.us-east-2.amazonaws.com | HTTPS |
| | | polly-fips.us-east-2.api.aws | HTTPS |
| | | polly-fips.us-east-2.amazonaws.com | HTTPS |
| | | polly.us-east-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | polly.us-east-1.amazonaws.com | HTTPS |
| | | polly-fips.us-east-1.api.aws | HTTPS |
| | | polly-fips.us-east-1.amazonaws.com | HTTPS |
| | | polly.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | polly.us-west-1.amazonaws.com | HTTPS |
| | | polly-fips.us-west-1.api.aws | HTTPS |
| | | polly-fips.us-west-1.amazonaws.com | HTTPS |
| | | polly.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | polly.us-west-2.amazonaws.com | HTTPS |
| | | polly-fips.us-west-2.api.aws | HTTPS |
| | | polly-fips.us-west-2.amazonaws.com | HTTPS |
| | | polly.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | polly.af-south-1.amazonaws.com | HTTPS |
| | | polly.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | polly.ap-east-1.amazonaws.com | HTTPS |
| | | polly.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | polly.ap-southeast-5.amazonaws.com | HTTPS |
| | | polly.ap-southeast-5.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | polly.ap-south-1.amazonaws.com<br>polly.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | polly.ap-northeast-3.amazonaws.com<br>polly.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | polly.ap-northeast-2.amazonaws.com<br>polly.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | polly.ap-southeast-1.amazonaws.com<br>polly.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | polly.ap-southeast-2.amazonaws.com<br>polly.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | polly.ap-southeast-7.amazonaws.com<br>polly.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | polly.ap-northeast-1.amazonaws.com<br>polly.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | polly.ca-central-1.amazonaws.com | HTTPS | |
| | | polly-fips.ca-central-1.api.aws | HTTPS | |
| | | polly-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | polly.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | polly.eu-central-1.amazonaws.com | HTTPS | |
| | | polly.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | polly.eu-west-1.amazonaws.com | HTTPS | |
| | | polly.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | polly.eu-west-2.amazonaws.com | HTTPS | |
| | | polly.eu-west-2.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | polly.eu-west-3.amazonaws.com | HTTPS | |
| | | polly.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | polly.eu-south-2.amazonaws.com | HTTPS | |
| | | polly.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | polly.eu-north-1.amazonaws.com | HTTPS | |
| | | polly.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | polly.eu-central-2.amazonaws.com | HTTPS | |
| | | polly.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | polly.me-south-1.amazonaws.com | HTTPS | |
| | | polly.me-south-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | polly.sa-east-1.amazonaws.com | HTTPS | |
| | | polly.sa-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | polly.us-gov-west-1.amazonaws.com | HTTPS | |
| | | polly-fips.us-gov-west-1.api.aws | HTTPS | |
| | | polly-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | polly.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Default StartSpeechSynthesisTask billed characters count | Each supported Region: 100,000 | Yes | Maximum size of the StartSpeechSynthesisTask input text in billed characters. SSML tags are not counted as billed characters. Applies to all engines. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default StartSpeechSynthesisTask lexicon count | Each supported Region: 5 | No | Maximum number of lexicons that can be used with StartSpeechSynthesisTask operation. Applies to all engines. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default StartSpeechSynthesisTask total characters limit | Each supported Region: 200,000 | Yes | Maximum size of the StartSpeechSynthesisTask input text in characters, including SSML tags and whitespace. Applies to all engines. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default SynthesizeSpeech billed character count | Each supported Region: 3,000 | Yes | Maximum size of the SynthesizeSpeech input text in billed characters. SSML tags are not counted as billed characters. Applies to all engines. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default SynthesizeSpeech lexicon count | Each supported Region: 5 | No | Maximum number of lexicons that can be used with SynthesizeSpeech operation. Applies to all engines. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default SynthesizeSpeech total character count | Each supported Region: 6,000 | Yes | Maximum size of the SynthesizeSpeech input text in characters, including SSML tags and whitespace. Applies to all engines. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default lexicon count | Each supported Region: 100 | No | The maximum number of lexicons you can have in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default lexicon size | Each supported Region: 40,000 | No | The maximum size of a single lexicon in characters. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default rate of GetSpeechSynthesisTask and ListSpeechSynthesisTasks requests | Each supported Region: 10 | Yes | The maximum number of GetSpeechSynthesisTask and ListSpeechSynthesisTasks requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default rate of StartSpeechSynthesisTask (generative) requests | Each supported Region: 1 | Yes | The maximum number of StartSpeechSynthesisTask (generative engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default rate of StartSpeechSynthesisTask (long-form) requests | Each supported Region: 1 | Yes | The maximum number of StartSpeechSynthesisTask (long-form engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default rate of StartSpeechSynthesisTask (neural) requests | Each supported Region: 10 | Yes | The maximum number of StartSpeechSynthesisTask (neural engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default rate of StartSpeechSynthesisTask (standard) requests | Each supported Region: 10 | Yes | The maximum number of StartSpeechSynthesisTask (standard engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default rate of SynthesizeSpeech (generative) requests | Each supported Region: 8 | Yes | The maximum number of SynthesizeSpeech (generative engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default rate of SynthesizeSpeech (long-form) requests | Each supported Region: 8 | Yes | The maximum number of SynthesizeSpeech (long-form engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default rate of SynthesizeSpeech (neural) requests | Each supported Region: 8 | Yes | The maximum number of SynthesizeSpeech (neural engine) requests per second that you can send in this account in the current region. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Default rate of SynthesizeSpeech (standard) requests | Each supported Region: 80 | Yes | The maximum number of SynthesizeSpeech (standard engine) requests per second that you can send in this account in the current region. This also includes DescribeVoices requests. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |
| Default rate of lexicon management requests | Each supported Region: 5 | Yes | The maximum number of lexicon management requests per second that you can send in this account in the current region. This limit applies to following operations combined: ListLexicon, GetLexicon, PutLexicon, DeleteLexicon. Please note that only the default limit is shown. If a higher limit has been granted to you, it wont be visible here. |

For more information, see Quotas in the *Amazon Polly Developer Guide*.

# AWS Private Certificate Authority endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## AWS Private CA endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | acm-pca.us-east-2.amazonaws.com | https |
| | | acm-pca-fips.us-east-2.amazonaws.com | HTTPS |
| | | acm-pca-fips.us-east-2.api.aws | HTTPS |
| | | acm-pca.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | acm-pca.us-east-1.amazonaws.com | https |
| | | acm-pca-fips.us-east-1.amazonaws.com | HTTPS |
| | | acm-pca-fips.us-east-1.api.aws | HTTPS |
| | | acm-pca.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | acm-pca.us-west-1.amazonaws.com | https |
| | | acm-pca-fips.us-west-1.amazonaws.com | HTTPS |
| | | acm-pca-fips.us-west-1.api.aws | HTTPS |
| | | acm-pca.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | acm-pca.us-west-2.amazonaws.com<br>acm-pca-fips.us-west-2.amazonaws.com<br>acm-pca-fips.us-west-2.api.aws<br>acm-pca.us-west-2.api.aws | https<br>HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | acm-pca.af-south-1.amazonaws.com<br>acm-pca.af-south-1.api.aws | https<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | acm-pca.ap-east-1.amazonaws.com<br>acm-pca.ap-east-1.api.aws | https<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | acm-pca.ap-south-2.amazonaws.com<br>acm-pca.ap-south-2.api.aws | https<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | acm-pca.ap-southeast-3.amazonaws.com<br>acm-pca.ap-southeast-3.api.aws | https<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | acm-pca.ap-southeast-5.amazonaws.com<br>acm-pca.ap-southeast-5.api.aws | https<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | acm-pca.ap-southeast-4.amazonaws.com<br>acm-pca.ap-southeast-4.api.aws | https<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | acm-pca.ap-south-1.amazonaws.com<br>acm-pca.ap-south-1.api.aws | https<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | acm-pca.ap-southeast-6.amazonaws.com<br>acm-pca.ap-southeast-6.api.aws | https<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | acm-pca.ap-northeast-3.amazonaws.com<br>acm-pca.ap-northeast-3.api.aws | https<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | acm-pca.ap-northeast-2.amazonaws.com<br>acm-pca.ap-northeast-2.api.aws | https<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | acm-pca.ap-southeast-1.amazonaws.com<br>acm-pca.ap-southeast-1.api.aws | https<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | acm-pca.ap-southeast-2.amazonaws.com<br>acm-pca.ap-southeast-2.api.aws | https<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | acm-pca.ap-east-2.amazonaws.com<br>acm-pca.ap-east-2.api.aws | https<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | acm-pca.ap-southeast-7.amazonaws.com<br>acm-pca.ap-southeast-7.api.aws | https<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | acm-pca.ap-northeast-1.amazonaws.com | https |
| | | acm-pca.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | acm-pca.ca-central-1.amazonaws.com | https |
| | | acm-pca-fips.ca-central-1.amazonaws.com | HTTPS |
| | | acm-pca-fips.ca-central-1.api.aws | HTTPS |
| | | acm-pca.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | acm-pca.ca-west-1.amazonaws.com | https |
| | | acm-pca-fips.ca-west-1.amazonaws.com | HTTPS |
| | | acm-pca-fips.ca-west-1.api.aws | HTTPS |
| | | acm-pca.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | acm-pca.eu-central-1.amazonaws.com | https |
| | | acm-pca.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | acm-pca.eu-west-1.amazonaws.com | https |
| | | acm-pca.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | acm-pca.eu-west-2.amazonaws.com | https |
| | | acm-pca.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | acm-pca.eu-south-1.amazonaws.com | https |
| | | acm-pca.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | acm-pca.eu-west-3.amazonaws.com | https |
| | | acm-pca.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | acm-pca.eu-south-2.amazonaws.com<br>acm-pca.eu-south-2.api.aws | https<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | acm-pca.eu-north-1.amazonaws.com<br>acm-pca.eu-north-1.api.aws | https<br>HTTPS | |
| Europe (Zurich) | eu-central-2 | acm-pca.eu-central-2.amazonaws.com<br>acm-pca.eu-central-2.api.aws | https<br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | acm-pca.il-central-1.amazonaws.com<br>acm-pca.il-central-1.api.aws | https<br>HTTPS | |
| Mexico (Central) | mx-central-1 | acm-pca.mx-central-1.amazonaws.com<br>acm-pca.mx-central-1.api.aws | https<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | acm-pca.me-south-1.amazonaws.com<br>acm-pca.me-south-1.api.aws | https<br>HTTPS | |
| Middle East (UAE) | me-central-1 | acm-pca.me-central-1.amazonaws.com<br>acm-pca.me-central-1.api.aws | https<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | acm-pca.sa-east-1.amazonaws.com<br>acm-pca.sa-east-1.api.aws | https<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | acm-pca.us-gov-east-1.amazonaws.com<br>acm-pca.us-gov-east-1.api.aws | https<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | acm-pca.us-gov-west-1.amazonaws.com<br>acm-pca.us-gov-west-1.api.aws | https<br>HTTPS | |

## AWS Private CA Connector for Active Directory endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | pca-connector-ad.us-east-2.amazonaws.com<br>pca-connector-ad.us-east-2.api.aws | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | pca-connector-ad.us-east-1.amazonaws.com<br>pca-connector-ad.us-east-1.api.aws | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | pca-connector-ad.us-west-1.amazonaws.com<br>pca-connector-ad.us-west-1.api.aws | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | pca-connector-ad.us-west-2.amazonaws.com<br>pca-connector-ad.us-west-2.api.aws | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | pca-connector-ad.af-south-1.amazonaws.com<br>pca-connector-ad.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific | ap-east-1 | pca-connector-ad.ap-east-1.amazonaws.com<br>pca-connector-ad.ap-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | pca-connector-ad.ap-south-2.amazonaws.com<br>pca-connector-ad.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | pca-connector-ad.ap-southeast-3.amazonaws.com<br>pca-connector-ad.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | pca-connector-ad.ap-southeast-5.amazonaws.com<br>pca-connector-ad.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | pca-connector-ad.ap-southeast-4.amazonaws.com<br>pca-connector-ad.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | pca-connector-ad.ap-south-1.amazonaws.com<br>pca-connector-ad.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | pca-connector-ad.ap-southeast-6.amazonaws.com<br>pca-connector-ad.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | pca-connector-ad.ap-northeast-3.amazonaws.com<br>pca-connector-ad.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | pca-connector-ad.ap-northeast-2.amazonaws.com<br>pca-connector-ad.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | pca-connector-ad.ap-southeast-1.amazonaws.com<br>pca-connector-ad.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | pca-connector-ad.ap-southeast-2.amazonaws.com<br>pca-connector-ad.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | pca-connector-ad.ap-east-2.amazonaws.com<br>pca-connector-ad.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | pca-connector-ad.ap-southeast-7.amazonaws.com<br>pca-connector-ad.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | pca-connector-ad.ap-northeast-1.amazonaws.com<br>pca-connector-ad.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | pca-connector-ad.ca-central-1.amazonaws.com<br>pca-connector-ad.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | pca-connector-ad.ca-west-1.amazonaws.com<br>pca-connector-ad.ca-west-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | pca-connector-ad.eu-central-1.amazonaws.com<br>pca-connector-ad.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | pca-connector-ad.eu-west-1.amazonaws.com<br>pca-connector-ad.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | pca-connector-ad.eu-west-2.amazonaws.com<br>pca-connector-ad.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | pca-connector-ad.eu-south-1.amazonaws.com<br>pca-connector-ad.eu-south-1.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | pca-connector-ad.eu-west-3.amazonaws.com<br>pca-connector-ad.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | pca-connector-ad.eu-south-2.amazonaws.com<br>pca-connector-ad.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | pca-connector-ad.eu-north-1.amazonaws.com<br>pca-connector-ad.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | pca-connector-ad.eu-central-2.amazonaws.com<br>pca-connector-ad.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | pca-connector-ad.il-central-1.amazonaws.com<br>pca-connector-ad.il-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | pca-connector-ad.mx-central-1.amazonaws.com<br>pca-connector-ad.mx-central-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | pca-connector-ad.me-south-1.amazonaws.com<br>pca-connector-ad.me-south-1.api.aws | HTTPS<br>HTTPS | |
| Middle East (UAE) | me-central-1 | pca-connector-ad.me-central-1.amazonaws.com<br>pca-connector-ad.me-central-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | pca-connector-ad.sa-east-1.amazonaws.com<br>pca-connector-ad.sa-east-1.api.aws | HTTPS<br>HTTPS | |

## AWS Private CA Connector for SCEP endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | pca-connector-scep.us-east-2.amazonaws.com<br>pca-connector-scep.us-east-2.api.aws | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | pca-connector-scep.us-east-1.amazonaws.com<br>pca-connector-scep.us-east-1.api.aws | HTTPS<br>HTTPS | |
| US West (N. | us-west-1 | pca-connector-scep.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| California) | | pca-connector-scep.us-west-1.api.aws | | |
| US West (Oregon) | us-west-2 | pca-connector-scep.us-west-2.amazonaws.com<br><br>pca-connector-scep.us-west-2.api.aws | HTTPS<br><br>HTTPS | |
| Africa (Cape Town) | af-south-1 | pca-connector-scep.af-south-1.amazonaws.com<br><br>pca-connector-scep.af-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | pca-connector-scep.ap-east-1.amazonaws.com<br><br>pca-connector-scep.ap-east-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | pca-connector-scep.ap-south-2.amazonaws.com<br><br>pca-connector-scep.ap-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | pca-connector-scep.ap-southeast-3.amazonaws.com<br><br>pca-connector-scep.ap-southeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | pca-connector-scep.ap-southeast-5.amazonaws.com<br><br>pca-connector-scep.ap-southeast-5.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | pca-connector-scep.ap-southeast-4.amazonaws.com<br><br>pca-connector-scep.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | pca-connector-scep.ap-south-1.amazonaws.com<br><br>pca-connector-scep.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | pca-connector-scep.ap-southeast-6.amazonaws.com<br><br>pca-connector-scep.ap-southeast-6.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | pca-connector-scep.ap-northeast-3.amazonaws.com<br><br>pca-connector-scep.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | pca-connector-scep.ap-northeast-2.amazonaws.com<br><br>pca-connector-scep.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | pca-connector-scep.ap-southeast-1.amazonaws.com<br><br>pca-connector-scep.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | pca-connector-scep.ap-southeast-2.amazonaws.com<br><br>pca-connector-scep.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | pca-connector-scep.ap-east-2.amazonaws.com<br><br>pca-connector-scep.ap-east-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | pca-connector-scep.ap-southeast-7.amazonaws.com<br><br>pca-connector-scep.ap-southeast-7.api.aws | HTTPS<br><br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | pca-connector-scep.ap-northeast-1.amazonaws.com<br><br>pca-connector-scep.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS |
| Canada (Central) | ca-central-1 | pca-connector-scep.ca-central-1.amazonaws.com<br><br>pca-connector-scep.ca-central-1.api.aws | HTTPS<br><br>HTTPS |
| Canada West (Calgary) | ca-west-1 | pca-connector-scep.ca-west-1.amazonaws.com<br><br>pca-connector-scep.ca-west-1.api.aws | HTTPS<br><br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | pca-connector-scep.eu-central-1.amazonaws.com<br><br>pca-connector-scep.eu-central-1.api.aws | HTTPS<br><br>HTTPS |
| Europe (Ireland) | eu-west-1 | pca-connector-scep.eu-west-1.amazonaws.com<br><br>pca-connector-scep.eu-west-1.api.aws | HTTPS<br><br>HTTPS |
| Europe (London) | eu-west-2 | pca-connector-scep.eu-west-2.amazonaws.com<br><br>pca-connector-scep.eu-west-2.api.aws | HTTPS<br><br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | pca-connector-scep.eu-south-1.amazonaws.com<br><br>pca-connector-scep.eu-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Paris) | eu-west-3 | pca-connector-scep.eu-west-3.amazonaws.com<br><br>pca-connector-scep.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Spain) | eu-south-2 | pca-connector-scep.eu-south-2.amazonaws.com<br><br>pca-connector-scep.eu-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | pca-connector-scep.eu-north-1.amazonaws.com<br><br>pca-connector-scep.eu-north-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Zurich) | eu-central-2 | pca-connector-scep.eu-central-2.amazonaws.com<br><br>pca-connector-scep.eu-central-2.api.aws | HTTPS<br><br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | pca-connector-scep.il-central-1.amazonaws.com<br><br>pca-connector-scep.il-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Mexico (Central) | mx-central-1 | pca-connector-scep.mx-central-1.amazonaws.com<br><br>pca-connector-scep.mx-central-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | pca-connector-scep.me-south-1.amazonaws.com<br><br>pca-connector-scep.me-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (UAE) | me-central-1 | pca-connector-scep.me-central-1.amazonaws.com<br><br>pca-connector-scep.me-central-1.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | pca-connector-scep.sa-east-1.amazonaws.com<br><br>pca-connector-scep.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## AWS Private CA quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of private certificate authorities | Each supported Region: 200 | Yes | The maximum number of private certificate authorities (CAs) that can be created in an account in the current Region. |
| Number of private certificates for each CA | Each supported Region: 100,000,000 | No | The maximum number of private certificates that a certificate authority (CA) can issue in an account in the current Region. This quota applies when no revocation is enabled or revocation is enabled |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | using partitioned CRL and/or OCSP. When complete CRL is enabled, the "Number of private certificates for each CA using complete CRLs" limit is applied. |
| Number of private certificates for each CA using complete CRLs | Each supported Region: 1,000,000 | No | The maximum number of private certificates that a certificate authority (CA) can issue in an account in the current Region. This quota applies to CAs using complete CRLs. |
| Number of revoked private certificates for each CA with complete CRLs enabled | Each supported Region: 1,000,000 | No | The maximum number of revoked, unexpired private certificates in an account in the current Region. This quota applies only to certificates issued when the CA had complete CRLs enabled. |
| Rate of CreateCertificateAuthority requests | Each supported Region: 1 | No | The maximum number of CreateCertificateAuthority requests that can be performed in an account in the current Region each second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateCertificateAuthorityAuditReport requests | Each supported Region: 1 | No | The maximum number of CreateCertificateAuthorityAuditReport requests that can be performed in an account in the current Region each second. |
| Rate of CreatePermission requests | Each supported Region: 1 | No | The maximum number of CreatePermission requests that can be performed in an account in the current Region each second. |
| Rate of DeleteCertificateAuthority requests | Each supported Region: 10 | No | The maximum number of DeleteCertificateAuthority requests that can be performed in an account in the current Region each second. |
| Rate of DeletePermission requests | Each supported Region: 1 | No | The maximum number of DeletePermission requests that can be performed in an account in the current Region each second. |
| Rate of DeletePolicy requests | Each supported Region: 5 | No | The maximum number of DeletePolicy requests that can be performed in an account in the current Region each second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeCertificateAuthority requests | Each supported Region: 20 | No | The maximum number of DescribeCertificateAuthority requests that can be performed in an account in the current Region each second. |
| Rate of DescribeCertificateAuthorityAuditReport requests | Each supported Region: 20 | No | The maximum number of DescribeCertificateAuthorityAuditReport requests that can be performed in an account in the current Region each second. |
| Rate of GetCertificate requests | Each supported Region: 75 | Yes | The maximum number of GetCertificate requests that can be performed in an account in the current Region each second. |
| Rate of GetCertificateAuthorityCertificate requests | Each supported Region: 20 | No | The maximum number of GetCertificateAuthorityCertificate requests that can be performed in an account in the current Region each second. |
| Rate of GetCertificateAuthorityCsr requests | Each supported Region: 10 | No | The maximum number of GetCertificateAuthorityCsr requests that can be performed in an account in the current Region each second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetPolicy requests | Each supported Region: 5 | No | The maximum number of GetPolicy requests that can be performed in an account in the current Region each second. |
| Rate of ImportCertificateAuthorityCertificate requests | Each supported Region: 10 | No | The maximum number of ImportCertificateAuthorityCertificate requests that can be performed in an account in the current Region each second. |
| Rate of IssueCertificate requests | Each supported Region: 25 | Yes | The maximum number of IssueCertificate requests that can be performed in an account in the current Region each second. |
| Rate of ListCertificateAuthorities requests | Each supported Region: 20 | No | The maximum number of ListCertificateAuthorities requests that can be performed in an account in the current Region each second. |
| Rate of ListPermissions requests | Each supported Region: 5 | No | The maximum number of ListPermissions requests that can be performed in an account in the current Region each second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListTags requests | Each supported Region: 20 | No | The maximum number of ListTags requests that can be performed in an account in the current Region each second. |
| Rate of PutPolicy requests | Each supported Region: 5 | No | The maximum number of PutPolicy requests that can be performed in an account in the current Region each second. |
| Rate of RestoreCertificateAuthority requests | Each supported Region: 20 | No | The maximum number of RestoreCertificate Authority requests that can be performed in an account in the current Region each second. |
| Rate of RevokeCertificate requests | Each supported Region: 20 | No | The maximum number of RevokeCertificate requests that can be performed in an account in the current Region each second. |
| Rate of TagCertificateAuthority requests | Each supported Region: 10 | No | The maximum number of TagCertificateAuthority requests that can be performed in an account in the current Region each second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UntagCertificateAuthority requests | Each supported Region: 10 | No | The maximum number of UntagCertificateAuthority requests that can be performed in an account in the current Region each second. |
| Rate of UpdateCertificateAuthority requests | Each supported Region: 10 | No | The maximum number of UpdateCertificateAuthority requests that can be performed in an account in the current Region each second. |

## AWS Private CA Connector for Active Directory quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of connectors | Each supported Region: 10 | Yes | The maximum number of connectors that you can create in this account in the current region. |
| Number of group access control entries per template | Each supported Region: 100 | Yes | The maximum number of group access control entries per template that you can create in this account in the current Region. |
| Number of templates per connector | Each supported Region: 100 | Yes | The maximum number of templates per connector |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | that you can create in this account in the current Region. |
| Rate of CreateConnector requests | Each supported Region: 2 | Yes | The maximum number of CreateConnector requests that you can perform in this account in the current region per second. |
| Rate of CreateDirectoryRegistration requests | Each supported Region: 10 | Yes | The maximum number of CreateDirectoryRegistration requests that you can perform in this account in the current region per second. |
| Rate of CreateServicePrincipalName requests | Each supported Region: 10 | Yes | The maximum number of CreateServicePrincipalName requests that you can perform in this account in the current region per second. |
| Rate of CreateTemplate requests | Each supported Region: 10 | Yes | The maximum number of CreateTemplate requests that you can perform in this account in the current region per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateTemplateGroupAccessControlEntry requests | Each supported Region: 10 | Yes | The maximum number of CreateTemplateGroupAccessControlEntry requests that you can perform in this account in the current region per second. |
| Rate of DeleteConnector requests | Each supported Region: 2 | Yes | The maximum number of DeleteConnector requests that you can perform in this account in the current region per second. |
| Rate of DeleteDirectoryRegistration requests | Each supported Region: 5 | Yes | The maximum number of DeleteDirectoryRegistration requests that you can perform in this account in the current region per second. |
| Rate of DeleteServicePrincipalName requests | Each supported Region: 5 | Yes | The maximum number of DeleteServicePrincipalName requests that you can perform in this account in the current region per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteTemplate requests | Each supported Region: 5 | Yes | The maximum number of DeleteTemplate requests that you can perform in this account in the current region per second. |
| Rate of DeleteTemplateGroupAccessControlEntry requests | Each supported Region: 5 | Yes | The maximum number of DeleteTemplateGroupAccessControlEntry requests that you can perform in this account in the current region per second. |
| Rate of GetConnector requests | Each supported Region: 50 | Yes | The maximum number of GetConnector requests that you can perform in this account in the current region per second. |
| Rate of GetDirectoryRegistration requests | Each supported Region: 50 | Yes | The maximum number of GetDirectoryRegistration requests that you can perform in this account in the current region per second. |
| Rate of GetPolicies requests | Each supported Region: 10 | Yes | The maximum number of GetPolicies requests that you can perform in this account in the current region per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of GetServicePrincipalName requests | Each supported Region: 50 | Yes | The maximum number of GetServicePrincipalName requests that you can perform in this account in the current region per second. |
| Rate of GetTemplate requests | Each supported Region: 50 | Yes | The maximum number of GetTemplate requests that you can perform in this account in the current region per second. |
| Rate of GetTemplateGroupAccessControlEntry requests | Each supported Region: 50 | Yes | The maximum number of GetTemplateGroupAccessControlEntry requests that you can perform in this account in the current region per second. |
| Rate of ListConnectors requests | Each supported Region: 50 | Yes | The maximum number of ListConnectors requests that you can perform in this account in the current region per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListDirectoryRegistrations requests | Each supported Region: 50 | Yes | The maximum number of ListDirectoryRegistrations requests that you can perform in this account in the current region per second. |
| Rate of ListServicePrincipalNames requests | Each supported Region: 50 | Yes | The maximum number of ListServicePrincipalNames requests that you can perform in this account in the current region per second. |
| Rate of ListTagsForResource requests | Each supported Region: 50 | Yes | The maximum number of ListTagsForResource requests that you can perform in this account in the current region per second. |
| Rate of ListTemplateGroupAccessControlEntries requests | Each supported Region: 50 | Yes | The maximum number of ListTemplateGroupAccessControlEntries requests that you can perform in this account in the current region per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListTemplates requests | Each supported Region: 50 | Yes | The maximum number of ListTemplates requests that you can perform in this account in the current region per second. |
| Rate of RequestSecurityToken requests | Each supported Region: 10 | Yes | The maximum number of RequestSecurityToken requests that you can perform in this account in the current region per second. |
| Rate of TagResource requests | Each supported Region: 50 | Yes | The maximum number of TagResource requests that you can perform in this account in the current region per second. |
| Rate of UntagResource requests | Each supported Region: 50 | Yes | The maximum number of UntagResource requests that you can perform in this account in the current region per second. |
| Rate of UpdateTemplate requests | Each supported Region: 10 | Yes | The maximum number of UpdateTemplate requests that you can perform in this account in the current region per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateTemplateGroupAccessControlEntry requests | Each supported Region: 10 | Yes | The maximum number of UpdateTemplateGroupAccessControlEntry requests that you can perform in this account in the current region per second. |

## AWS Private CA Connector for SCEP quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of challenges per connector | Each supported Region: 5 | Yes | The maximum number of challenges per connector that you can create in this account in the current Region. |
| Number of connectors | Each supported Region: 20 | Yes | The maximum number of connectors that you can create in this account in the current Region. |
| Rate of CreateChallenge requests | Each supported Region: 5 | Yes | The maximum number of CreateChallenge requests per second that you can perform in this account in the current Region. |
| Rate of CreateConnector requests | Each supported Region: 5 | Yes | The maximum number of CreateConnector requests per second that |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | you can perform in this account in the current Region. |
| Rate of DeleteChallenge requests | Each supported Region: 5 | Yes | The maximum number of DeleteChallenge requests per second that you can perform in this account in the current Region. |
| Rate of DeleteConnector requests | Each supported Region: 5 | Yes | The maximum number of DeleteConnector requests per second that you can perform in this account in the current Region. |
| Rate of GetCACaps requests | Each supported Region: 20 | Yes | The maximum number of GetCACaps requests per second that your connector can perform in the current Region. |
| Rate of GetCACert requests | Each supported Region: 20 | Yes | The maximum number of GetCACert requests per second that your connector can perform in the current Region. |
| Rate of GetChallengeMetadata requests | Each supported Region: 50 | Yes | The maximum number of GetChallengeMetadata requests per second that you can perform in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of GetChallengePassword requests | Each supported Region: 50 | Yes | The maximum number of GetChallengePassword requests per second that you can perform in this account in the current Region. |
| Rate of GetConnector requests | Each supported Region: 50 | Yes | The maximum number of GetConnector requests per second that you can perform in this account in the current Region. |
| Rate of ListChallengeMetadata requests | Each supported Region: 50 | Yes | The maximum number of ListChallengeMetadata requests per second that you can perform in this account in the current Region. |
| Rate of ListConnectors requests | Each supported Region: 50 | Yes | The maximum number of ListConnectors requests per second that you can perform in this account in the current Region. |
| Rate of ListTagsForResource requests | Each supported Region: 50 | Yes | The maximum number of ListTagsForResource requests per second that you can perform in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of PKCSReq requests | Each supported Region: 10 | Yes | The maximum number of PKCSReq requests per second that your connector can perform in the current Region. |
| Rate of TagResource requests | Each supported Region: 50 | Yes | The maximum number of TagResource requests per second that you can perform in this account in the current Region. |
| Rate of UntagResource requests | Each supported Region: 50 | Yes | The maximum number of UntagResource requests per second that you can perform in this account in the current Region. |

# Amazon Managed Service for Prometheus endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

Amazon Managed Service for Prometheus includes control plane endpoints (to perform workspace management tasks) and data plane endpoints (to work with Prometheus-compatible data in a workspace instance). Control plane endpoints start with `aps.*`, and dataplane endpoints start with `aps-workspaces.*`. Endpoints that end in `.amazonaws.com` support IPv4, and endpoints that end in `.api.aws` support both IPv4 and IPv6.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | aps.us-east-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.us-east-2.amazonaws.com | HTTPS | |
| | | aps-workspaces-fips.us-east-2.amazonaws.com | HTTPS | |
| | | aps-workspaces-fips.us-east-2.api.aws | HTTPS | |
| | | aps-workspaces.us-east-2.api.aws | HTTPS | |
| | | aps-fips.us-east-2.amazonaws.com | HTTPS | |
| | | aps.us-east-2.api.aws | HTTPS | |
| | | aps-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | aps.us-east-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.us-east-1.amazonaws.com | HTTPS | |
| | | aps-workspaces-fips.us-east-1.amazonaws.com | HTTPS | |
| | | aps-workspaces-fips.us-east-1.api.aws | HTTPS | |
| | | aps-workspaces.us-east-1.api.aws | HTTPS | |
| | | aps-fips.us-east-1.amazonaws.com | HTTPS | |
| | | | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | aps.us-east-1.api.aws | HTTPS |
| | | aps-fips.us-east-1.api.aws | |
| US West (N. California) | us-west-1 | aps.us-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces.us-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-west-1.api.aws | HTTPS |
| | | aps-workspaces.us-west-1.api.aws | HTTPS |
| | | aps-fips.us-west-1.amazonaws.com | HTTPS |
| | | aps.us-west-1.api.aws | HTTPS |
| | | aps-fips.us-west-1.api.aws | |
| US West (Oregon) | us-west-2 | aps.us-west-2.amazonaws.com | HTTPS |
| | | aps-workspaces.us-west-2.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-west-2.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-west-2.api.aws | HTTPS |
| | | aps-workspaces.us-west-2.api.aws | HTTPS |
| | | aps-fips.us-west-2.amazonaws.com | HTTPS |
| | | aps.us-west-2.api.aws | HTTPS |
| | | aps-fips.us-west-2.api.aws | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | aps.af-south-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.af-south-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.af-south-1.api.aws | HTTPS | |
| | | aps.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | aps.ap-east-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.ap-east-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.ap-east-1.api.aws | HTTPS | |
| | | aps.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | aps.ap-south-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.ap-south-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.ap-south-2.api.aws | HTTPS | |
| | | aps.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | aps.ap-southeast-3.amazonaws.com | HTTPS | |
| | | aps-workspaces.ap-southeast-3.amazonaws.com | HTTPS | |
| | | aps-workspaces.ap-southeast-3.api.aws | HTTPS | |
| | | aps.ap-southeast-3.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | aps.ap-southeast-5.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-5.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-5.api.aws | HTTPS |
| | | aps.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | aps.ap-southeast-4.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-4.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-4.api.aws | HTTPS |
| | | aps.ap-southeast-4.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | aps.ap-south-1.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-south-1.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-south-1.api.aws | HTTPS |
| | | aps.ap-south-1.api.aws | HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | aps.ap-southeast-6.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-6.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-6.api.aws | HTTPS |
| | | aps.ap-southeast-6.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | aps.ap-northeast-3.amazonaws.com<br>aps-workspaces.ap-northeast-3.amazonaws.com<br>aps-workspaces.ap-northeast-3.api.aws<br>aps.ap-northeast-3.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | aps.ap-northeast-2.amazonaws.com<br>aps-workspaces.ap-northeast-2.amazonaws.com<br>aps-workspaces.ap-northeast-2.api.aws<br>aps.ap-northeast-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | aps.ap-southeast-1.amazonaws.com<br>aps-workspaces.ap-southeast-1.amazonaws.com<br>aps-workspaces.ap-southeast-1.api.aws<br>aps.ap-southeast-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | aps.ap-southeast-2.amazonaws.com<br>aps-workspaces.ap-southeast-2.amazonaws.com<br>aps-workspaces.ap-southeast-2.api.aws<br>aps.ap-southeast-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | aps.ap-east-2.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-east-2.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-east-2.api.aws | HTTPS |
| | | aps.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | aps.ap-southeast-7.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-7.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-southeast-7.api.aws | HTTPS |
| | | aps.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | aps.ap-northeast-1.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-northeast-1.amazonaws.com | HTTPS |
| | | aps-workspaces.ap-northeast-1.api.aws | HTTPS |
| | | aps.ap-northeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | aps.ca-central-1.amazonaws.com<br>aps-workspaces.ca-central-1.amazonaws.com<br>aps-workspaces-fips.ca-central-1.amazonaws.com<br>aps-workspaces-fips.ca-central-1.api.aws<br>aps-workspaces.ca-central-1.api.aws<br>aps-fips.ca-central-1.amazonaws.com<br>aps.ca-central-1.api.aws<br>aps-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | aps.ca-west-1.amazonaws.com<br>aps-workspaces.ca-west-1.amazonaws.com<br>aps-workspaces-fips.ca-west-1.amazonaws.com<br>aps-workspaces-fips.ca-west-1.api.aws<br>aps-workspaces.ca-west-1.api.aws<br>aps-fips.ca-west-1.amazonaws.com<br>aps.ca-west-1.api.aws<br>aps-fips.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | aps.eu-central-1.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-central-1.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-central-1.api.aws | HTTPS |
| | | aps.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | aps.eu-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-west-1.api.aws | HTTPS |
| | | aps.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | aps.eu-west-2.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-west-2.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-west-2.api.aws | HTTPS |
| | | aps.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | aps.eu-south-1.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-south-1.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-south-1.api.aws | HTTPS |
| | | aps.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | aps.eu-west-3.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-west-3.amazonaws.com | HTTPS |
| | | aps-workspaces.eu-west-3.api.aws | HTTPS |
| | | aps.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | aps.eu-south-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.eu-south-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.eu-south-2.api.aws | HTTPS | |
| | | aps.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | aps.eu-north-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.eu-north-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.eu-north-1.api.aws | HTTPS | |
| | | aps.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | aps.eu-central-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.eu-central-2.amazonaws.com | HTTPS | |
| | | aps-workspaces.eu-central-2.api.aws | HTTPS | |
| | | aps.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | aps.il-central-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.il-central-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.il-central-1.api.aws | HTTPS | |
| | | aps.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | aps.mx-central-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.mx-central-1.amazonaws.com | HTTPS | |
| | | aps-workspaces.mx-central-1.api.aws | HTTPS | |
| | | aps.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | aps.me-south-1.amazonaws.com | HTTPS |
| | | aps-workspaces.me-south-1.amazonaws.com | HTTPS |
| | | aps-workspaces.me-south-1.api.aws | HTTPS |
| | | aps.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | aps.me-central-1.amazonaws.com | HTTPS |
| | | aps-workspaces.me-central-1.amazonaws.com | HTTPS |
| | | aps-workspaces.me-central-1.api.aws | HTTPS |
| | | aps.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | aps.sa-east-1.amazonaws.com | HTTPS |
| | | aps-workspaces.sa-east-1.amazonaws.com | HTTPS |
| | | aps-workspaces.sa-east-1.api.aws | HTTPS |
| | | aps.sa-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | aps.us-gov-east-1.amazonaws.com | HTTPS |
| | | aps-workspaces.us-gov-east-1.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-gov-east-1.api.aws | HTTPS |
| | | aps-workspaces.us-gov-east-1.api.aws | HTTPS |
| | | aps-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | aps.us-gov-east-1.api.aws | HTTPS |
| | | aps-fips.us-gov-east-1.api.aws | |
| AWS GovCloud (US-West) | us-gov-west-1 | aps.us-gov-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces.us-gov-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | aps-workspaces-fips.us-gov-west-1.api.aws | HTTPS |
| | | aps-workspaces.us-gov-west-1.api.aws | HTTPS |
| | | aps-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | aps.us-gov-west-1.api.aws | HTTPS |
| | | aps-fips.us-gov-west-1.api.aws | |

# Service quotas

Amazon Managed Service for Prometheus has the following quotas. Amazon Managed Service for Prometheus vends CloudWatch usage metrics to monitor Prometheus resource usage. Using the CloudWatch usage metrics alarm feature, you can monitor Prometheus resources and usage to prevent limit errors.

As your projects and workspaces grow, the most common quotas that you may need to monitor or request an increase for are: **Active series per workspace**, **Ingestion rate per workspace**, and **Ingestion burst size per workspace**.

For all adjustable quotas, you can request a quota increase by selecting the link in the **Adjustable** column, or by requesting a quota increase.

> **Note**
>
> Unless otherwise noted, these quotas are per workspace.

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Active metrics with metadata per workspace | Each supported Region: 20,000 | No | The number of unique active metrics with metadata per workspace. Note: If the limit is reached, metric sample is recorded, but metadata over the limit is dropped. |
| Active series per workspace | Each supported Region: 50,000,000 | Yes | The number of unique active series per workspace (up to a maximum of 1.5 billion). A series is active if a sample has been reported in the past 2 hours. Capacity from 2 M |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | to 50 M is automatically adjusted based on the last 30 min of usage. |
| Alert aggregation group size in alert manager definition file | Each supported Region: 1,000 | Yes | The maximum size of an alert aggregation group in alert manager definition file. Each label value combination of group_by would create an aggregation group. |
| Alert manager definition file size | Each supported Region: 1,000,000 | No | The maximum size of an alert manager definition file, in bytes. |
| Alert payload size in Alert Manager | Each supported Region: 20,000,000 | No | The maximum alert payload size of all Alert Manager alerts per workspace, in bytes. Alert size is dependent on labels and annotations. |
| Alerts in Alert Manager | Each supported Region: 1,000 | Yes | The maximum number of concurrent Alert Manager alerts per workspace. |
| HA tracker clusters | Each supported Region: 500 | No | The maximum number of clusters that HA tracker will keep track of for ingested samples per workspace. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Ingestion rate per workspace | Each supported Region: 1,666,666 | Yes | Metric sample ingestion rate per workspace per second. The limit is automatically adjusted to be 1/30 of the active series per workspace limit, up to 1,666,666. |
| Inhibition rules in alert manager definition file | Each supported Region: 100 | Yes | The maximum number of inhibition rules in alert manager definition file. |
| Label size | Each supported Region: 7 | No | The maximum combined size of all labels and label values accepted for a series, in kilobytes. |
| LabelSet limits per workspace | Each supported Region: 100 | Yes | The maximum number of labelset limits that can be created per workspace . |
| Labels per metric series | Each supported Region: 150 | Yes | Number of labels per metric series. |
| Max Native Histogram bucket count | Each supported Region: 200 | No | The maximum total number of buckets allowed in a Native Histogram sample. The resolution of a histogram with more buckets will be reduced until the number of buckets is within the limit. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Metadata length | Each supported Region: 1 | No | The maximum length accepted for metric metadata, in kilobytes. Metadata refers to Metric Name, Type, Unit and Help Text. |
| Metadata per metric | Each supported Region: 10 | No | The number of metadata per metric. Note: If the limit is reached, metric sample is recorded, but metadata over the limit is dropped. |
| Native Histogram active series per workspace | Each supported Region: 5,000,000 | Yes | The number of unique Native Histogram active series per workspace . A series is active if a sample has been reported in the past 2 hours. The limit is automatically adjusted to be 10-percent of the workspaces active series limit, or the default, whichever is greater. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Native Histogram ingestion rate per workspace | Each supported Region: 166,666 | Yes | Native Histogram sample ingestion rate per workspace per second. The limit is automatically adjusted to be 10-percent of the workspaces ingestion rate limit, or the default, whichever is greater. |
| Nodes in alert manager routing tree | Each supported Region: 100 | Yes | The maximum number of nodes in the alert manager routing tree. |
| Number of API operations per region in transactions per second | Each supported Region: 10 | Yes | The maximum number of API operations per second per region for all Amazon Managed Service for Prometheus APIs, including workspace CRUD APIs, tagging APIs, rule groups namespace CRUD APIs, and alert manager definition CRUD APIs. |
| Number of GetSeries, GetLabels and GetMetricMetadata API operations per workspace in transactions per second | Each supported Region: 10 | No | The maximum number of GetSeries, GetLabels and GetMetricMetadata Prometheus-compatible API operations per second per workspace. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of QueryMetrics API operations per workspace in transactions per second | Each supported Region: 300 | No | The maximum number of QueryMetrics Prometheus-compatible API operations per second per workspace. |
| Number of RemoteWrite API operations per workspace in transactions per second | Each supported Region: 3,000 | No | The maximum number of RemoteWrite Prometheus-compatible API operations per second per workspace. |
| Number of other Prometheus-compatible API operations per workspace in transactions per second | Each supported Region: 100 | No | The maximum number of API operations per second per workspace for all other Prometheus-compatible APIs including ListAlerts, ListRules, etc. |
| Out of order ingestion rate per workspace | Each supported Region: 83,333 | Yes | Out of order samples ingestion rate per workspace per second. Unless overwritten, the limit is automatically adjusted to be 5-percent of the ingestion rate per workspace limit. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Out of order time window per workspace | Each supported Region: 600 | Yes | The maximum time window for out of order samples per workspace, in seconds. The default out-of-order time window when a new workspace is created is 60 seconds, and it can be configured up to a maximum of 600 seconds. |
| Query bytes for instant queries | Each supported Region: 5 | No | The maximum bytes that can be scanned by a single instant query, in gigabytes. |
| Query bytes for range queries | Each supported Region: 5 | No | The maximum bytes that can be scanned per 24-hour interval in a single range query, in gigabytes. |
| Query samples | Each supported Region: 50,000,000 | No | The maximum number of samples that can be scanned per 24-hour interval in a single range query, or a single instant query. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Query series fetched | Each supported Region: 12,000,000 | No | The maximum number of series that can be scanned per 24-hour interval in a single range query, or a single instant query. |
| Query time range in days | Each supported Region: 95 | No | The maximum time range of QueryMetrics, GetSeries, and GetLabels APIs. |
| Request size | Each supported Region: 1 | No | The maximum request size for ingestion or query, in megabytes. |
| Rule evaluation interval | Each supported Region: 30 | Yes | The minimum rule evaluation interval of a rule group per workspace, in seconds. |
| Rule group namespace definition file size | Each supported Region: 1,000,000 | No | The maximum size of a rule group namespace definition file, in bytes. |
| Rules per workspace | Each supported Region: 2,000 | Yes | The maximum number of rules per workspace. |
| Silences per workspace | Each supported Region: 1,000 | Yes | Maximum number of silences, including expired, active, and pending silences, per workspace. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Templates in alert manager definition file | Each supported Region: 100 | Yes | The maximum number of templates in the alert manager definition file. |
| Workspaces per region per account | Each supported Region: 25 | Yes | The maximum number of workspaces per region. |

### Additional limits on ingested data

Amazon Managed Service for Prometheus also has the following additional requirements for data ingested into the workspace. These are not adjustable.

- Metric samples older than 1 hour are refused from being ingested.
- Every sample and metadata must have a metric name.

## AWS Proton endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

### Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | proton.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | proton.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | proton.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | proton.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | proton.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | proton.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | proton.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | proton.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | proton.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | proton.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | proton.eu-west-2.amazonaws.com | HTTPS | |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Components per account | Each supported Region: 1,000 | Yes | Maximum number of components per account |
| Environment account connections per environment account | Each supported Region: 1,000 | Yes | Maximum number of environment account connections per environment account |
| Environments per account | Each supported Region: 1,000 | Yes | Maximum number of environments per account |
| Service instances per service | Each supported Region: 20 | Yes | Maximum number of service instances per service |
| Services per account | Each supported Region: 1,000 | Yes | Maximum number of services per account |
| Template versions per template | Each supported Region: 1,000 | Yes | Maximum number of template versions registered per template |
| Templates per account | Each supported Region: 1,000 | Yes | Maximum number of registered templates per account, service and environment templates combined |

For more information, see AWS Proton quotas in the *AWS Proton Administrator Guide*.

# Amazon Q Business endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service regions and endpoints

The following table shows the AWS Regions and endpoints currently supported by Amazon Q Business.

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | qbusiness.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | qbusiness.us-west-2.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | qbusiness.eu-west-1.api.aws<br><br>qbusiness-websocket.eu-west-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | qbusiness.ap-southeast-2.api.aws<br><br>qbusiness-websocket.ap-southeast-2.api.aws | HTTPS |

**Note**

The Europe (Ireland) and the Asia Pacific (Sydney) regions don't currently support all features available in the US regions, such as Q App, Q Actions, and Audio/Video files. While these features will become available soon, Amazon Q Business customers in this region can do the following:

- Get answers to questions submitted to the enterprise retrieval augmented generation system.
- Generate content through Amazon Q Business assistant
- Access capabilities such as embedded images in files.
- Perform tabular search on small tables.
- Ingest data from scanned PDFs.
- Answer questions from data in scanned PDFs.
- Respond to queries to LLM knowledge.

## Service quotas

The following table shows the quotas that are related to Amazon Q Business for your AWS account.

| Name | Description | Default | Adjustable |
|---|---|---|---|
| Applications | Maximum number of applications per account | 50 | No |
| Data sources | Maximum number of data sources per application | 50 | No |
| Data accessors | Maximum number of data accessors per application | 10 | No |

| Name | Description | Default | Adjustable |
| --- | --- | --- | --- |
| Plugins | Maximum number of plugins per applicati on | 25 | No |
| Actions | Maximum number of actions per plugin | 20 | No |

# Amazon Q Developer endpoints and quotas

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. The following are the service quotas for this service.

## Amazon Q Developer Transform service quotas

| Feature | Quota (per user) |
| --- | --- |
| .NET monthly lines of code (Q Dev Pro Subscription) | 1,000,000 lines of code per supported Region |
| Concurrent .NET jobs with the Microsoft Visual Studio Extension (Q Dev Pro Subscriber) | 10 jobs per supported Region |
| .NET monthly lines of code with the Microsoft Visual Studio Extension when used with BuilderIDs | 1,000 lines of code per supported Region |
| Concurrent .NET jobs with the Microsoft Visual Studio Extension (BuilderID) | 1 job per supported Region |

## Amazon Q Developer Pro service quotas

The features included with Amazon Q Pro subscriptions are subject to the usage quotas outlined on the Amazon Q Developer pricing page and in the following tables. These quotas ensure fair usage and optimal performance for all subscribers.

For more information about pricing, see the Amazon Q Developer pricing page.

### Quotas for Amazon Q Developer Pro (in IDEs and CLI)

| Feature | Quota (per user) |
| --- | --- |
| Amazon Q Developer Agent for code transformation | 4,000 lines of code per month, pooled at the account level. Extra lines of code available. For more information, see Amazon Q Developer Pro pricing. |
| Agentic requests (Q&A chat, agentic coding) | 1,000 lines of code per month |

### Quotas for Amazon Q Developer Pro (in the AWS Management Console)

| Feature | Quota (per user) |
| --- | --- |
| Generative SQL | 1,000 queries per month |
| Analyze network reachability | 20 requests per day |

### Quotas for Amazon Q Developer Pro (in Amazon CodeCatalyst)

| Feature | Quota (per user) |
| --- | --- |
| Amazon Q Developer Agent for software development | 30 per month |
| Pull request summaries | 20 per month |

## Systems Manager Quick Setup endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ssm-quicksetup.us-east-2.amazonaws.com | HTTPS |
| | | ssm-quicksetup-fips.us-east-2.api.aws | HTTPS |
| | | ssm-quicksetup-fips.us-east-2.amazonaws.com | HTTPS |
| | | ssm-quicksetup.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | ssm-quicksetup.us-east-1.amazonaws.com | HTTPS |
| | | ssm-quicksetup-fips.us-east-1.api.aws | HTTPS |
| | | ssm-quicksetup-fips.us-east-1.amazonaws.com | HTTPS |
| | | ssm-quicksetup.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ssm-quicksetup.us-west-1.amazonaws.com | HTTPS |
| | | ssm-quicksetup-fips.us-west-1.api.aws | HTTPS |
| | | ssm-quicksetup-fips.us-west-1.amazonaws.com | HTTPS |
| | | ssm-quicksetup.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ssm-quicksetup.us-west-2.amazonaws.com | HTTPS |
| | | ssm-quicksetup-fips.us-west-2.api.aws | HTTPS |
| | | | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | ssm-quicksetup-fips.us-west-2.amazonaws.com<br>ssm-quicksetup.us-west-2.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ssm-quicksetup.ap-south-1.amazonaws.com<br>ssm-quicksetup.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ssm-quicksetup.ap-northeast-2.amazonaws.com<br>ssm-quicksetup.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ssm-quicksetup.ap-southeast-1.amazonaws.com<br>ssm-quicksetup.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ssm-quicksetup.ap-southeast-2.amazonaws.com<br>ssm-quicksetup.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ssm-quicksetup.ap-northeast-1.amazonaws.com<br>ssm-quicksetup.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | ssm-quicksetup.ca-central-1.amazonaws.com<br>ssm-quicksetup-fips.ca-central-1.api.aws<br>ssm-quicksetup-fips.ca-central-1.amazonaws.com<br>ssm-quicksetup.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | ssm-quicksetup.eu-central-1.amazonaws.com<br>ssm-quicksetup.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | ssm-quicksetup.eu-west-1.amazonaws.com<br>ssm-quicksetup.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | ssm-quicksetup.eu-west-2.amazonaws.com<br>ssm-quicksetup.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | ssm-quicksetup.eu-west-3.amazonaws.com<br>ssm-quicksetup.eu-west-3.api.aws | HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | ssm-quicksetup.eu-north-1.amazonaws.com<br>ssm-quicksetup.eu-north-1.api.aws | HTTPS<br>HTTPS | |
| South America (São Paulo) | sa-east-1 | ssm-quicksetup.sa-east-1.amazonaws.com<br>ssm-quicksetup.sa-east-1.api.aws | HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of configuration managers | Each supported Region: 100 | No | The maximum number of Quick Setup configuration managers you can create in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateConfigurationManager API requests | Each supported Region: 1 | No | The maximum number of CreateConfigurationManager API requests that you can make per second in this account in the current Region. |
| Rate of DeleteConfigurationManager API requests | Each supported Region: 3 | No | The maximum number of DeleteConfigurationManager API requests that you can make per second in this account in the current Region. |
| Rate of GetConfiguration API requests | Each supported Region: 5 | No | The maximum number of GetConfiguration API requests that you can make per second in this account in the current Region. |
| Rate of GetConfigurationManager API requests | Each supported Region: 5 | No | The maximum number of GetConfigurationManager API requests that you can make per second in this account in the current Region. |
| Rate of GetServiceSettings API requests | Each supported Region: 5 | No | The maximum number of GetServiceSettings API requests that you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListConfigurationManagers API requests | Each supported Region: 5 | No | The maximum number of ListConfigurationManagers API requests that you can make per second in this account in the current Region. |
| Rate of ListConfigurations API requests | Each supported Region: 5 | No | The maximum number of ListConfigurations API requests that you can make per second in this account in the current Region. |
| Rate of ListQuickSetupTypes API requests | Each supported Region: 5 | No | The maximum number of ListQuickSetupTypes API requests that you can make per second in this account in the current Region. |
| Rate of ListTagsForResource API requests | Each supported Region: 5 | No | The maximum number of ListTagsForResource API requests that you can make per second in this account in the current Region. |
| Rate of TagResource API requests | Each supported Region: 2 | No | The maximum number of TagResource API requests that you can make per second in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of UntagResource API requests | Each supported Region: 2 | No | The maximum number of UntagResource API requests that you can make per second in this account in the current Region. |
| Rate of UpdateConfigurationDefinition API requests | Each supported Region: 1 | No | The maximum number of UpdateConfigurationDefinition API requests that you can make per second in this account in the current Region. |
| Rate of UpdateConfigurationManager API requests | Each supported Region: 1 | No | The maximum number of UpdateConfigurationManager API requests that you can make per second in this account in the current Region. |
| Rate of UpdateServiceSettings API requests | Each supported Region: 5 | No | The maximum number of UpdateServiceSettings API requests that you can make per second in this account in the current Region. |

# Amazon Quick Sight

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## QuickSight

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | quicksight.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | quicksight.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | quicksight.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | quicksight.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | quicksight.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | quicksight.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | quicksight.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | quicksight.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | quicksight.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | quicksight.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | quicksight.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | quicksight.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | quicksight.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | quicksight.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | quicksight.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | quicksight.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | quicksight.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | quicksight.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | quicksight.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | quicksight.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | quicksight.il-central-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | quicksight.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | quicksight.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | quicksight.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | quicksight.us-gov-west-1.amazonaws.com | HTTPS | |

## QuickSight Websites

| Region Name | Region | Endpoint |
|---|---|---|
| US East (Ohio) | us-east-2 | https://us-east-2.quicksight.aws.amazon.com |
| US East (N. Virginia) | us-east-1 | https://us-east-1.quicksight.aws.amazon.com |
| US West (Oregon) | us-west-2 | https://us-west-2.quicksight.aws.amazon.com |
| Canada (Central) | ca-central-1 | https://ca-central-1.quicksight.aws.amazon.com |
| Asia Pacific (Mumbai) | ap-south-1 | https://ap-south-1.quicksight.aws.amazon.com |
| Asia Pacific (Singapore) | ap-southeast-1 | https://ap-southeast-1.quicksight.aws.amazon.com |
| Asia Pacific (Sydney) | ap-southeast-2 | https://ap-southeast-2.quicksight.aws.amazon.com |
| Asia Pacific (Tokyo) | ap-northeast-1 | https://ap-northeast-1.quicksight.aws.amazon.com |
| Asia Pacific (Seoul) | ap-northeast-2 | https://ap-northeast-2.quicksight.aws.amazon.com |
| Europe (Frankfurt) | eu-central-1 | https://eu-central-1.quicksight.aws.amazon.com |
| Europe (Ireland) | eu-west-1 | https://eu-west-1.quicksight.aws.amazon.com |
| Europe (London) | eu-west-2 | https://eu-west-2.quicksight.aws.amazon.com |
| Europe (Paris) | eu-west-3 | https://eu-west-3.quicksight.aws.amazon.com |
| Europe (Stockholm) | eu-north-1 | https://eu-north-1.quicksight.aws.amazon.com |

| Region Name | Region | Endpoint |
|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | https://quicksight.us-gov-east-1.amazonaws.com |
| AWS GovCloud (US-West) | us-gov-west-1 | https://quicksight.us-gov-west-1.amazonaws.com |
| Israel (Tel Aviv) | il-central-1 | https://il-central-1.quicksight.aws.amazon.com/ |
| Middle East (UAE) | me-central-1 | https://me-central-1.quicksight.aws.amazon.com/ |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| API_CREATE-INGESTION: Calls per 24 hour period from Enterprise edition | Each supported Region: 32 | No | The maximum number of calls to the createIngestion API function in a floating 24-hour window. The time period is measured starting 24 hours before the current date and time. This maximum applies to AWS accounts that use Amazon QuickSight Enterprise edition. |
| API_CREATE-INGESTION: Calls per 24 hour period from Standard edition | Each supported Region: 8 | No | The maximum number of calls to the createIngestion API function in a floating 24-hour window. The time period is measured starting 24 hours before the |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | current date and time. This maximum applies to AWS accounts that use Amazon QuickSight Standard edition. |
| Amazon Quick Automate Maximum Automation Groups | Each supported Region: 200 | Yes | Maximum number of Automation Groups allowed per account in Amazon Quick Automate |
| Amazon Quick Automate Maximum Automations | Each supported Region: 1,000 | Yes | Maximum number of Automations allowed across all Automation Groups in Amazon Quick Automate |
| Amazon Quick Automate Maximum Concurrent Executions per Trigger | Each supported Region: 20 | No | Maximum number of concurrent executions per trigger in Amazon Quick Automate |
| Amazon Quick Automate Maximum Concurrent Tests and Deployed Executions | Each supported Region: 100 | Yes | Maximum number of concurrent tests and executions per account in Amazon Quick Automate |
| Amazon Quick Automate Maximum Credentials | Each supported Region: 200 | Yes | Maximum number of credentials allowed across all Automation Groups in Amazon Quick Automate |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Amazon Quick Automate Maximum Nesting Depth | Each supported Region: 6 | No | Maximum nesting depth allowed in an automation in Amazon Quick Automate |
| Amazon Quick Automate Maximum Triggers per Automation | Each supported Region: 30 | No | Maximum number of triggers allowed per automation in Amazon Quick Automate |
| Amazon Quick Automate Maximum Workflow Execution Time | Each supported Region: 48 | No | Maximum workflow execution time in hours in Amazon Quick Automate |
| Calculated field expression length | Each supported Region: 250,000 | No | The maximum number of characters that you can use in an expression for a calculated field. |
| Conversation history retention in days | us-east-1: 365<br>eu-west-1: 365<br>Each of the other supported Regions: 90 | No | The number of days conversation history is retained before deletion. |
| Custom action name length | Each supported Region: 256 | No | The maximum number of characters that you can use in naming a custom action. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom actions per visual | Each supported Region: 10 | No | The maximum number of custom actions that you can configure for each visual in an analysis. |
| Custom agent artifact retention in days | Each supported Region: 365 | No | The number of days custom agent artifacts are retained before deletion. |
| Data Prep: Fields per dataset | Each supported Region: 2,000 | No | The maximum number of fields that a dataset can contain. File imports and query result sets can contain more than 2,000 columns. However, you must edit the dataset settings and manually exclude fields until there are less than 2,000 selected or included. |
| Display items per sheet control | Each supported Region: 1,000 | No | The maximum number of distinct items that a sheet control can display. |
| Email aliases per group for email reports | Each supported Region: 5,000 | No | The maximum number of members in any group that QuickSight sends email reports to. If you try to send reports to larger groups, the report fails. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| In-folder conversation history retention in days | Each supported Region: 750 | No | The number of days conversation history in a folder is retained before deletion. |
| Maximum active sessions per user | Each supported Region: 3 | No | The maximum number of concurrent active sessions a user can have. |
| Maximum applicable groups per approval policy | Each supported Region: 5 | Yes | The maximum number of group ARNs that can be specified in an approval policys ApplicableTo configuration (the groups the policy applies to). |
| Maximum approver groups per approval policy | Each supported Region: 5 | Yes | The maximum number of approver group ARNs that can be specified in an approval policys ApprovalGroups configuration (the groups that can approve requests). |
| Maximum artifact file size | Each supported Region: 500 Megabytes | No | The maximum size in megabytes (MBs) of a single artifact stored in personal space. |
| Maximum concurrent runs per user | Each supported Region: 30 | No | The maximum number of active conversations with concurrent runs a user can have. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum concurrent tool executions | Each supported Region: 20 | No | The maximum number of tool executions that can run concurrently within a conversation. |
| Maximum conversation folders | Each supported Region: 20 | No | The maximum number of conversation folders a user can create. |
| Maximum custom agent files size | Each supported Region: 1 Gigabytes | No | The maximum total size in gigabytes (GBs) of files associated with a custom agent. |
| Maximum custom agent instruction length | Each supported Region: 50,000 | No | The maximum number of characters allowed in a custom agent instruction. |
| Maximum file attachment size | Each supported Region: 50 Megabytes | No | The maximum size in megabytes (MBs) of a single file attachment. |
| Maximum installed and draft custom agents | Each supported Region: 100 | No | The maximum combined number of installed and draft custom agents a user can have. |
| Maximum installed connectors | Each supported Region: 100 | No | The maximum number of connectors (A2A and MCP) a user can have installed. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum memory and knowledge graph size | Each supported Region: 2 Gigabytes | No | The maximum combined size in gigabytes (GBs) of memory and knowledge graph data per user. |
| Maximum number of approval policies per account | Each supported Region: 500 | Yes | The maximum number of approval policies that can be created in a single account across all asset types. |
| Maximum number of approval policies per asset type | Each supported Region: 100 | Yes | The maximum number of approval policies per asset type (for example, KNOWLEDGE_BASE) within a single account. |
| Maximum number of characters per specified Control values | Each supported Region: 200,000 | No | The maximum number of characters used in the entries that you type in to display inside sheet controls. An example is the values specified for a dropdown. This doesnt apply to values created from a dataset. |
| Maximum number of reference documents per chat agent | Each supported Region: 10 | No | The maximum number of reference documents that can be attached to a chat agent. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of resources per chat agent | Each supported Region: 20 | No | The maximum number of resources that can be associated with a chat agent. |
| Maximum number of resources per space | Each supported Region: 100 | No | The maximum number of resources that can be linked to a space. |
| Maximum number of uploaded files per space | Each supported Region: 10,000 | No | The maximum number of files that can be uploaded to a space. |
| Maximum pinned conversations | Each supported Region: 100 | No | The maximum number of conversations a user can pin. |
| Maximum reference file size | Each supported Region: 50 Megabytes | No | The maximum size in megabytes (MBs) of a single reference file. |
| Maximum scheduled task instruction length | Each supported Region: 50,000 | No | The maximum number of characters allowed in a scheduled task instruction. |
| Maximum scheduled tasks per user | Each supported Region: 20 | No | The maximum number of scheduled tasks a user can create across all triggers. |
| Maximum skill size | Each supported Region: 6 Megabytes | No | The maximum size in megabytes (MBs) of a single skill bundle. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum skills per user | Each supported Region: 1,000 | No | The maximum number of skills a user can create. |
| Maximum total artifact storage per user | Each supported Region: 51,200 Megabytes | No | The maximum total size in megabytes (MBs) of all artifacts stored in a users personal space. |
| Maximum total data capacity of an index | Each supported Region: 60,000 Megabytes | No | The maximum total data capacity in megabytes (MBs) of an index. |
| Maximum total size of attached reference documents per chat agent | Each supported Region: 50 Megabytes | No | The maximum size in megabytes (MBs) of all reference documents attached to a chat agent. |
| Maximum total size of uploaded files per space | Each supported Region: 1 Gigabytes | No | The maximum size in gigabytes (GBs) of all files uploaded to a space. |
| Minimum schedule interval in minutes | Each supported Region: 15 | No | The minimum interval, in minutes, allowed between scheduled task runs. |
| Pinned and in-folder conversation history retention in days | Each supported Region: 750 | No | The number of days conversation history that is both pinned and in a folder is retained before deletion. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Pinned conversation history retention in days | Each supported Region: 750 | No | The number of days pinned conversation history is retained before deletion. |
| Query timeout for visuals | Each supported Region: 120 Seconds | No | The maximum amount of time that QuickSight waits for a database to finish sending data. This applies to queries initiated by visuals. |
| Scheduled task history retention in days | us-east-1: 365<br>eu-west-1: 365<br>Each of the other supported Regions: 90 | No | The number of days scheduled task run history is retained. |
| Session approval expiry in hours | Each supported Region: 12 | No | The number of hours before a session approval expires and must be renewed. |
| The maximum amount of time to wait for a dataset preview | Each supported Region: 45 Seconds | No | The maximum amount of time that QuickSight waits for a data preview to finish loading. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| URL action hyperlink length | Each supported Region: 2,048 | No | The maximum number of characters allowed in the hyperlink (URL) of a custom action thats defined as a URL action. This includes all variations of the link for the different parameters you include. |

# AWS Resource Access Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | ram.us-east-2.amazonaws.com | HTTPS | |
| | | ram-fips.us-east-2.api.aws | HTTPS | |
| | | ram.us-east-2.api.aws | HTTPS | |
| | | ram-fips.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ram.us-east-1.amazonaws.com | HTTPS |
| | | ram-fips.us-east-1.api.aws | HTTPS |
| | | ram-fips.us-east-1.amazonaws.com | HTTPS |
| | | ram.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ram.us-west-1.amazonaws.com | HTTPS |
| | | ram.us-west-1.api.aws | HTTPS |
| | | ram-fips.us-west-1.api.aws | HTTPS |
| | | ram-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | ram.us-west-2.amazonaws.com | HTTPS |
| | | ram-fips.us-west-2.api.aws | HTTPS |
| | | ram.us-west-2.api.aws | HTTPS |
| | | ram-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | ram.af-south-1.amazonaws.com | HTTPS |
| | | ram.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | ram.ap-east-1.amazonaws.com | HTTPS |
| | | ram.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | ram.ap-south-2.amazonaws.com | HTTPS |
| | | ram.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | ram.ap-southeast-3.amazonaws.com<br>ram.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | ram.ap-southeast-5.amazonaws.com<br>ram.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | ram.ap-southeast-4.amazonaws.com<br>ram.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ram.ap-south-1.amazonaws.com<br>ram.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | ram.ap-southeast-6.amazonaws.com<br>ram.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | ram.ap-northeast-3.amazonaws.com<br>ram.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ram.ap-northeast-2.amazonaws.com<br>ram.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ram.ap-southeast-1.amazonaws.com<br>ram.ap-southeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | ram.ap-southeast-2.amazonaws.com | HTTPS |
| | | ram.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | ram.ap-east-2.amazonaws.com | HTTPS |
| | | ram.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | ram.ap-southeast-7.amazonaws.com | HTTPS |
| | | ram.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ram.ap-northeast-1.amazonaws.com | HTTPS |
| | | ram.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | ram.ca-central-1.amazonaws.com | HTTPS |
| | | ram.ca-central-1.api.aws | HTTPS |
| | | ram-fips.ca-central-1.api.aws | HTTPS |
| | | ram-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | ram.ca-west-1.amazonaws.com | HTTPS |
| | | ram.ca-west-1.api.aws | HTTPS |
| | | ram-fips.ca-west-1.api.aws | HTTPS |
| | | ram-fips.ca-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ram.eu-central-1.amazonaws.com | HTTPS |
| | | ram.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | ram.eu-west-1.amazonaws.com | HTTPS | |
| | | ram.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | ram.eu-west-2.amazonaws.com | HTTPS | |
| | | ram.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | ram.eu-south-1.amazonaws.com | HTTPS | |
| | | ram.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | ram.eu-west-3.amazonaws.com | HTTPS | |
| | | ram.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | ram.eu-south-2.amazonaws.com | HTTPS | |
| | | ram.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | ram.eu-north-1.amazonaws.com | HTTPS | |
| | | ram.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | ram.eu-central-2.amazonaws.com | HTTPS | |
| | | ram.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | ram.il-central-1.amazonaws.com | HTTPS | |
| | | ram.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | ram.mx-central-1.amazonaws.com | HTTPS | |
| | | ram.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | ram.me-south-1.amazonaws.com<br>ram.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | ram.me-central-1.amazonaws.com<br>ram.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | ram.sa-east-1.amazonaws.com<br>ram.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | ram.us-gov-east-1.amazonaws.com<br>ram.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | ram.us-gov-west-1.amazonaws.com<br>ram.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of custom permissions | Each supported Region: 1,500 | Yes | The maximum number of custom permissions. |
| Number of custom permissions per resource type | Each supported Region: 10 | Yes | The maximum number of custom permissions that you can create per resource type. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of pending invitations | Each supported Region: 250 | Yes | The maximum number of pending invitations. |
| Number of principal associations | Each supported Region: 25,000 | Yes | The maximum number of principal associations. |
| Number of principal associations per resource share | Each supported Region: 5,000 | No | The maximum number of principals that you can specify in a resource share. |
| Number of resource associations | Each supported Region: 25,000 | Yes | The maximum number of resources associations. |
| Number of resource associations per resource share | Each supported Region: 5,000 | No | The maximum number of resources that you can include in a resource share. |
| Number of resource shares | Each supported Region: 25,000 | Yes | The maximum number of resource shares. |

**Notes**

- The quota for **Number of pending invitations** applies to only *sending* accounts who share with accounts that are ***not*** part of sender's AWS Organization.
- There is no quota for how many pending invitations a receiving account can have.
- Invitations are not used when sharing between accounts that are part of the same AWS Organization and resource sharing within that AWS Organization is turned on.

# Recycle Bin endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Topics**


- Service endpoints
- Service quotas


## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | rbin.us-east-2.amazonaws.com | HTTPS |
| | | rbin-fips.us-east-2.api.aws | HTTPS |
| | | rbin-fips.us-east-2.amazonaws.com | HTTPS |
| | | rbin.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | rbin.us-east-1.amazonaws.com | HTTPS |
| | | rbin-fips.us-east-1.api.aws | HTTPS |
| | | rbin-fips.us-east-1.amazonaws.com | HTTPS |
| | | rbin.us-east-1.api.aws | HTTPS |
| US West (N. | us-west-1 | rbin.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| California) | | rbin-fips.us-west-1.api.aws | HTTPS | |
| | | rbin-fips.us-west-1.amazonaws.com | HTTPS | |
| | | rbin.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | rbin.us-west-2.amazonaws.com | HTTPS | |
| | | rbin-fips.us-west-2.api.aws | HTTPS | |
| | | rbin-fips.us-west-2.amazonaws.com | HTTPS | |
| | | rbin.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | rbin.af-south-1.amazonaws.com | HTTPS | |
| | | rbin.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | rbin.ap-east-1.amazonaws.com | HTTPS | |
| | | rbin.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | rbin.ap-south-2.amazonaws.com | HTTPS | |
| | | rbin.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rbin.ap-southeast-3.amazonaws.com | HTTPS | |
| | | rbin.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rbin.ap-southeast-5.amazonaws.com | HTTPS | |
| | | rbin.ap-southeast-5.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | rbin.ap-southeast-4.amazonaws.com<br>rbin.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | rbin.ap-south-1.amazonaws.com<br>rbin.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | rbin.ap-southeast-6.amazonaws.com<br>rbin.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | rbin.ap-northeast-3.amazonaws.com<br>rbin.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | rbin.ap-northeast-2.amazonaws.com<br>rbin.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | rbin.ap-southeast-1.amazonaws.com<br>rbin.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | rbin.ap-southeast-2.amazonaws.com<br>rbin.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | rbin.ap-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | rbin.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | rbin.ap-northeast-1.amazonaws.com | HTTPS |
| | | rbin.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | rbin.ca-central-1.amazonaws.com | HTTPS |
| | | rbin-fips.ca-central-1.api.aws | HTTPS |
| | | rbin-fips.ca-central-1.amazonaws.com | HTTPS |
| | | rbin.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | rbin.ca-west-1.amazonaws.com | HTTPS |
| | | rbin-fips.ca-west-1.api.aws | HTTPS |
| | | rbin-fips.ca-west-1.amazonaws.com | HTTPS |
| | | rbin.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | rbin.eu-central-1.amazonaws.com | HTTPS |
| | | rbin.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | rbin.eu-west-1.amazonaws.com | HTTPS |
| | | rbin.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | rbin.eu-west-2.amazonaws.com | HTTPS |
| | | rbin.eu-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | rbin.eu-south-1.amazonaws.com | HTTPS | |
| | | rbin.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | rbin.eu-west-3.amazonaws.com | HTTPS | |
| | | rbin.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | rbin.eu-south-2.amazonaws.com | HTTPS | |
| | | rbin.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | rbin.eu-north-1.amazonaws.com | HTTPS | |
| | | rbin.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | rbin.eu-central-2.amazonaws.com | HTTPS | |
| | | rbin.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | rbin.il-central-1.amazonaws.com | HTTPS | |
| | | rbin.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | rbin.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | rbin.me-south-1.amazonaws.com | HTTPS | |
| | | rbin.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | rbin.me-central-1.amazonaws.com | HTTPS | |
| | | rbin.me-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| South America (São Paulo) | sa-east-1 | rbin.sa-east-1.amazonaws.com | HTTPS |
| | | rbin.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | rbin.us-gov-east-1.amazonaws.com | HTTPS |
| | | rbin-fips.us-gov-east-1.api.aws | HTTPS |
| | | rbin-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | rbin.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | rbin.us-gov-west-1.amazonaws.com | HTTPS |
| | | rbin-fips.us-gov-west-1.api.aws | HTTPS |
| | | rbin-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | rbin.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Quota | Default quota | Adjustable |
|---|---|---|
| Retention rules per Region | 250 | No |
| Tag key and value pairs per retention rule | 50 | No |

# Amazon Redshift endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Redshift API

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | redshift.us-east-2.amazonaws.com | HTTPS | |
| | | redshift-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | redshift.us-east-1.amazonaws.com | HTTPS | |
| | | redshift-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | redshift.us-west-1.amazonaws.com | HTTPS | |
| | | redshift-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | redshift.us-west-2.amazonaws.com | HTTPS | |
| | | redshift-fips.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | redshift.af-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | redshift.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | redshift.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | redshift.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | redshift.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | redshift.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | redshift.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | redshift.ap-southeast-6.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | redshift.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | redshift.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | redshift.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | redshift.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | redshift.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | redshift.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | redshift.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | redshift.ca-central-1.amazonaws.com<br>redshift-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | redshift.ca-west-1.amazonaws.com | HTTPS | |
| | | redshift-fips.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | redshift.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | redshift.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | redshift.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | redshift.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | redshift.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | redshift.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | redshift.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | redshift.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | redshift.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | redshift.mx-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | redshift.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | redshift.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | redshift.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | redshift.us-gov-east-1.amazonaws.com<br>redshift-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | redshift.us-gov-west-1.amazonaws.com<br>redshift-fips.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

**Redshift Serverless API**

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | redshift-serverless.us-east-2.amazonaws.com<br>redshift-serverless-fips.us-east-2.api.aws<br>redshift-serverless-fips.us-east-2.amazonaws.com<br>redshift-serverless.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | redshift-serverless.us-east-1.amazonaws.com<br>redshift-serverless-fips.us-east-1.api.aws<br>redshift-serverless-fips.us-east-1.amazonaws.com<br>redshift-serverless.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | redshift-serverless.us-west-1.amazonaws.com<br>redshift-serverless-fips.us-west-1.api.aws<br>redshift-serverless-fips.us-west-1.amazonaws.com<br>redshift-serverless.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | redshift-serverless.us-west-2.amazonaws.com<br>redshift-serverless-fips.us-west-2.api.aws<br>redshift-serverless-fips.us-west-2.amazonaws.com<br>redshift-serverless.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | redshift-serverless.af-south-1.amazonaws.com<br>redshift-serverless.af-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | redshift-serverless.ap-east-1.amazonaws.com<br>redshift-serverless.ap-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | redshift-serverless.ap-south-2.amazonaws.com<br>redshift-serverless.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | redshift-serverless.ap-southeast-3.amazonaws.com<br>redshift-serverless.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | redshift-serverless.ap-southeast-5.amazonaws.com<br>redshift-serverless.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | redshift-serverless.ap-southeast-4.amazonaws.com<br>redshift-serverless.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | redshift-serverless.ap-south-1.amazonaws.com<br>redshift-serverless.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | redshift-serverless.ap-southeast-6.amazonaws.com<br>redshift-serverless.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | redshift-serverless.ap-northeast-3.amazonaws.com<br>redshift-serverless.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | redshift-serverless.ap-northeast-2.amazonaws.com<br>redshift-serverless.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | redshift-serverless.ap-southeast-1.amazonaws.com<br>redshift-serverless.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | redshift-serverless.ap-southeast-2.amazonaws.com<br>redshift-serverless.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | redshift-serverless.ap-east-2.amazonaws.com<br>redshift-serverless.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | redshift-serverless.ap-southeast-7.amazonaws.com<br>redshift-serverless.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | redshift-serverless.ap-northeast-1.amazonaws.com<br>redshift-serverless.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | redshift-serverless.ca-central-1.amazonaws.com<br>redshift-serverless-fips.ca-central-1.api.aws<br>redshift-serverless-fips.ca-central-1.amazonaws.com<br>redshift-serverless.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | redshift-serverless.ca-west-1.amazonaws.com<br>redshift-serverless-fips.ca-west-1.api.aws<br>redshift-serverless-fips.ca-west-1.amazonaws.com<br>redshift-serverless.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | redshift-serverless.eu-central-1.amazonaws.com<br>redshift-serverless.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | redshift-serverless.eu-west-1.amazonaws.com<br>redshift-serverless.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | redshift-serverless.eu-west-2.amazonaws.com<br>redshift-serverless.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | redshift-serverless.eu-south-1.amazonaws.com<br>redshift-serverless.eu-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | redshift-serverless.eu-west-3.amazonaws.com<br>redshift-serverless.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | redshift-serverless.eu-south-2.amazonaws.com<br>redshift-serverless.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | redshift-serverless.eu-north-1.amazonaws.com<br>redshift-serverless.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | redshift-serverless.eu-central-2.amazonaws.com<br>redshift-serverless.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | redshift-serverless.il-central-1.amazonaws.com<br>redshift-serverless.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | redshift-serverless.mx-central-1.amazonaws.com<br>redshift-serverless.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | redshift-serverless.me-central-1.amazonaws.com<br>redshift-serverless.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | redshift-serverless.sa-east-1.amazonaws.com<br>redshift-serverless.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | redshift-serverless.us-gov-east-1.amazonaws.com<br>redshift-serverless-fips.us-gov-east-1.api.aws<br>redshift-serverless-fips.us-gov-east-1.amazonaws.com<br>redshift-serverless.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | redshift-serverless.us-gov-west-1.amazonaws.com<br>redshift-serverless-fips.us-gov-west-1.api.aws<br>redshift-serverless-fips.us-gov-west-1.amazonaws.com<br>redshift-serverless.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

**Redshift Data API**

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | redshift-data.us-east-2.amazonaws.com<br>redshift-data-fips.us-east-2.api.aws<br>redshift-data.us-east-2.api.aws<br>redshift-data-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | redshift-data.us-east-1.amazonaws.com<br>redshift-data-fips.us-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | redshift-data.us-east-1.api.aws | HTTPS |
| | | redshift-data-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | redshift-data.us-west-1.amazonaws.com | HTTPS |
| | | redshift-data-fips.us-west-1.api.aws | HTTPS |
| | | redshift-data.us-west-1.api.aws | HTTPS |
| | | redshift-data-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | redshift-data.us-west-2.amazonaws.com | HTTPS |
| | | redshift-data-fips.us-west-2.api.aws | HTTPS |
| | | redshift-data.us-west-2.api.aws | HTTPS |
| | | redshift-data-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | redshift-data.af-south-1.amazonaws.com | HTTPS |
| | | redshift-data.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | redshift-data.ap-east-1.amazonaws.com | HTTPS |
| | | redshift-data.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | redshift-data.ap-south-2.amazonaws.com | HTTPS |
| | | redshift-data.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | redshift-data.ap-southeast-3.amazonaws.com | HTTPS |
| | | redshift-data.ap-southeast-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | redshift-data.ap-southeast-5.amazonaws.com<br>redshift-data.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | redshift-data.ap-southeast-4.amazonaws.com<br>redshift-data.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | redshift-data.ap-south-1.amazonaws.com<br>redshift-data.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | redshift-data.ap-southeast-6.amazonaws.com<br>redshift-data.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | redshift-data.ap-northeast-3.amazonaws.com<br>redshift-data.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | redshift-data.ap-northeast-2.amazonaws.com<br>redshift-data.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | redshift-data.ap-southeast-1.amazonaws.com<br>redshift-data.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | redshift-data.ap-southeast-2.amazonaws.com<br>redshift-data.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | redshift-data.ap-east-2.amazonaws.com | HTTPS |
| | | redshift-data.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | redshift-data.ap-southeast-7.amazonaws.com | HTTPS |
| | | redshift-data.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | redshift-data.ap-northeast-1.amazonaws.com | HTTPS |
| | | redshift-data.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | redshift-data.ca-central-1.amazonaws.com | HTTPS |
| | | redshift-data-fips.ca-central-1.api.aws | HTTPS |
| | | redshift-data.ca-central-1.api.aws | HTTPS |
| | | redshift-data-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | redshift-data.ca-west-1.amazonaws.com | HTTPS |
| | | redshift-data-fips.ca-west-1.api.aws | HTTPS |
| | | redshift-data.ca-west-1.api.aws | HTTPS |
| | | redshift-data-fips.ca-west-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | redshift-data.eu-central-1.amazonaws.com | HTTPS |
| | | redshift-data.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | redshift-data.eu-west-1.amazonaws.com | HTTPS |
| | | redshift-data.eu-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | redshift-data.eu-west-2.amazonaws.com | HTTPS |
| | | redshift-data.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | redshift-data.eu-south-1.amazonaws.com | HTTPS |
| | | redshift-data.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | redshift-data.eu-west-3.amazonaws.com | HTTPS |
| | | redshift-data.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | redshift-data.eu-south-2.amazonaws.com | HTTPS |
| | | redshift-data.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | redshift-data.eu-north-1.amazonaws.com | HTTPS |
| | | redshift-data.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | redshift-data.eu-central-2.amazonaws.com | HTTPS |
| | | redshift-data.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | redshift-data.il-central-1.amazonaws.com | HTTPS |
| | | redshift-data.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | redshift-data.mx-central-1.amazonaws.com | HTTPS |
| | | redshift-data.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | redshift-data.me-south-1.amazonaws.com | HTTPS |
| | | redshift-data.me-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | redshift-data.me-central-1.amazonaws.com | HTTPS |
| | | redshift-data.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | redshift-data.sa-east-1.amazonaws.com | HTTPS |
| | | redshift-data.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | redshift-data.us-gov-east-1.amazonaws.com | HTTPS |
| | | redshift-data-fips.us-gov-east-1.api.aws | HTTPS |
| | | redshift-data.us-gov-east-1.api.aws | HTTPS |
| | | redshift-data-fips.us-gov-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | redshift-data.us-gov-west-1.amazonaws.com | HTTPS |
| | | redshift-data-fips.us-gov-west-1.api.aws | HTTPS |
| | | redshift-data.us-gov-west-1.api.aws | HTTPS |
| | | redshift-data-fips.us-gov-west-1.amazonaws.com | HTTPS |

**Redshift query editor v2**

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | api.sqlworkbench.us-east-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | api.sqlworkbench.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | api.sqlworkbench.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | api.sqlworkbench.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | api.sqlworkbench.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | api.sqlworkbench.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | api.sqlworkbench.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | api.sqlworkbench.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | api.sqlworkbench.ap-southeast-5.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | api.sqlworkbench.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | api.sqlworkbench.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | api.sqlworkbench.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | api.sqlworkbench.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | api.sqlworkbench.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | api.sqlworkbench.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | api.sqlworkbench.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | api.sqlworkbench.ap-southeast-7.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | api.sqlworkbench.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | api.sqlworkbench.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | api.sqlworkbench.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | api.sqlworkbench.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | api.sqlworkbench.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | api.sqlworkbench.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | api.sqlworkbench.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | api.sqlworkbench.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | api.sqlworkbench.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | api.sqlworkbench.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | api.sqlworkbench.eu-central-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | api.sqlworkbench.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | api.sqlworkbench.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | api.sqlworkbench.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | api.sqlworkbench.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | api.sqlworkbench.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | api.sqlworkbench.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | api.sqlworkbench.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

For information, see Quotas and limits in Amazon Redshift in the *Amazon Redshift Management Guide*.

# Amazon Rekognition endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Amazon Rekognition API operations (excluding streaming API operations) are available at the following regions and endpoints:

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | rekognition.us-east-2.amazonaws.com | HTTPS |
| | | rekognition.us-east-2.api.aws | HTTPS |
| | | rekognition-fips.us-east-2.amazonaws.com | HTTPS |
| | | rekognition-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | rekognition.us-east-1.amazonaws.com | HTTPS |
| | | rekognition-fips.us-east-1.amazonaws.com | HTTPS |
| | | rekognition.us-east-1.api.aws | HTTPS |
| | | rekognition-fips.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | rekognition.us-west-1.amazonaws.com | HTTPS |
| | | rekognition.us-west-1.api.aws | HTTPS |
| | | rekognition-fips.us-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | rekognition-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | rekognition.us-west-2.amazonaws.com | HTTPS |
| | | rekognition-fips.us-west-2.amazonaws.com | HTTPS |
| | | rekognition.us-west-2.api.aws | HTTPS |
| | | rekognition-fips.us-west-2.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | rekognition.ap-southeast-5.amazonaws.com | HTTPS |
| | | rekognition.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | rekognition.ap-south-1.amazonaws.com | HTTPS |
| | | rekognition.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | rekognition.ap-northeast-2.amazonaws.com | HTTPS |
| | | rekognition.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | rekognition.ap-southeast-1.amazonaws.com | HTTPS |
| | | rekognition.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | rekognition.ap-southeast-2.amazonaws.com | HTTPS |
| | | rekognition.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | rekognition.ap-southeast-7.amazonaws.com | HTTPS |
| | | rekognition.ap-southeast-7.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | rekognition.ap-northeast-1.amazonaws.com | HTTPS |
| | | rekognition.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | rekognition.ca-central-1.amazonaws.com | HTTPS |
| | | rekognition.ca-central-1.api.aws | HTTPS |
| | | rekognition-fips.ca-central-1.amazonaws.com | HTTPS |
| | | rekognition-fips.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | rekognition.eu-central-1.amazonaws.com | HTTPS |
| | | rekognition.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | rekognition.eu-west-1.amazonaws.com | HTTPS |
| | | rekognition.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | rekognition.eu-west-2.amazonaws.com | HTTPS |
| | | rekognition.eu-west-2.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | rekognition.eu-south-2.amazonaws.com | HTTPS |
| | | rekognition.eu-south-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | rekognition.il-central-1.amazonaws.com | HTTPS |
| | | rekognition.il-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | rekognition.sa-east-1.amazonaws.com | HTTPS |
| | | rekognition.sa-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | rekognition.us-gov-west-1.amazonaws.com | HTTPS |
| | | rekognition-fips.us-gov-west-1.api.aws | HTTPS |
| | | rekognition-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | rekognition.us-gov-west-1.api.aws | HTTPS |

## Amazon Rekognition Streaming Endpoints

The Amazon Rekognition streaming API operations are available at the following regions and endpoints:

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | streaming-rekognition.us-east-1.amazonaws.com | WSS |
| | | streaming-rekognition-fips.us-east-1.amazonaws.com | WSS |
| US West (Oregon) | us-west-2 | streaming-rekognition.us-west-2.amazonaws.com | WSS |
| | | streaming-rekognition-fips.us-west-2.amazonaws.com | WSS |
| Asia Pacific (Mumbai) | ap-south-1 | streaming-rekognition.ap-south-1.amazonaws.com | WSS |

| | | | |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | streaming-rekognition.ap-northeast-1.amazonaws.com | WSS |
| Europe (Ireland) | eu-west-1 | streaming-rekognition.eu-west-1.amazonaws.com | WSS |
| South America (São Paulo) | sa-east-1 | streaming-rekognition.sa-east-1.amazonaws.com | WSS |
| Asia Pacific (Malaysia) | ap-southeast-5 | streaming-rekognition.ap-southeast-5.amazonaws.com | WSS |
| Asia Pacific (Thailand) | ap-southeast-7 | streaming-rekognition.ap-southeast-7.amazonaws.com | WSS |

**Note**

Some regions only support certain Amazon Rekognition feature or operations. See the sections below for information on these differences.

The following are differences for certain Amazon Rekognition features and AWS Regions.

## Amazon Rekognition Video streaming API

The Amazon Rekognition Video streaming API is available in the following regions, depending on the specified Settings when creating a StreamProcessor.

Label Detection (ConnectedHome) API:

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)

- Asia Pacific (Mumbai)
- Europe (Ireland)

Face Search (FaceSearch) API:

- US East (N. Virginia)
- US West (Oregon)
- Asia Pacific (Tokyo)
- Europe (Frankfurt)
- Europe (Ireland)

## Amazon Rekognition Custom Labels

Amazon Rekognition Custom Labels is available in the following Regions only.

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- Europe (Ireland)
- Europe (London)
- Europe (Frankfurt)
- Asia Pacific (Mumbai)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Asia Pacific (Seoul)

## Canada (Central) Region

The Canada (Central) Region supports the following operations only.

- AssociateFaces
- CompareFaces

- CreateCollection
- DeleteCollection
- DeleteFaces
- DescribeCollection
- DetectFaces
- DisassociateFaces
- IndexFaces
- ListCollections
- ListFaces
- SearchFaces
- SearchFacesByImage
- SearchUsers
- SearchUsersByImage
- CreateUser
- DeleteUser
- ListUsers

> **Note**
>
> These operations are only available through use of the AWS CLI or SDK, as the Canada (Central) Region doesn't currently provide a console experience for these operations.

## Israel (Tel Aviv) Region

The Israel (Tel Aviv) Region supports only the following operations for the following features.

| Feature | Operations |
| --- | --- |
| Face detection | DetectFaces |
| Face comparison | CompareFaces |

| Feature | Operations |
| --- | --- |
| Face search | AssociateFaces, CreateCollection, CreateUser, DeleteCollection, DeleteFaces, DeleteUser, DescribeCollection, DisassociateFaces, IndexFaces, ListCollections, ListFaces, ListUsers, SearchFaces, SearchFacesByImage, SearchUsers, SearchUsersByImage, TagResource, UntagResource, ListTagsForResource |
| Label detection | DetectLabels |
| Moderation | DetectModerationLabels |
| Text detection | DetectText |

## Service quotas

The quotas listed on this page are defaults. You can request a quota increase for Amazon Rekognition using the AWS Support Center. To request a quota increase for a Amazon Rekognition Transactions Per Second (TPS) limit, follow the instructions at Default quotas in the *Amazon Rekognition Developer Guide*.

Quotas increases affect only the specific API operation for the Region in which you make the request. Other API operations and Regions are not affected.

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual Amazon Rekognition Image data plane operations:<br>• DetectLabels<br>• DetectModerationLabels<br>• DetectText<br>• | • US East (Ohio) Region – 5<br>• US East (N. Virginia) Region – 50<br>• US West (N. California) Region – 5 |

| Resource | Default |
| --- | --- |
| GetCelebrityInfo<br>• IndexFaces<br>• ListFaces<br>• RecognizeCelebrities<br>• SearchFaces<br>• SearchFacesByImage<br>• SearchUsers<br>• SearchUsersByImage | • US West (Oregon) Region – 50<br>• Asia Pacific (Mumbai) Region – 5<br>• Asia Pacific (Seoul) Region – 5<br>• Asia Pacific (Singapore) Region – 5<br>• Asia Pacific (Sydney) Region – 5<br>• Asia Pacific (Tokyo) Region – 5<br>• Canada (Central) – 5 (For supported operations, see Canada (Central) Region).<br>• Europe (Frankfurt) Region – 5<br>• Europe (Spain) Region – 5 |

| Resource | Default |
| --- | --- |
| | • South America (São Paulo) Region – 5<br>• Asia Pacific (Malaysia) Region – 5<br>• Asia Pacific (Thailand) Region – 5<br>• Europe (Ireland) Region – 50<br>• Europe (London) Region – 5<br>• Israel (Tel Aviv) Region – 5 (For supported operations, see Israel (Tel Aviv) Region)<br>• AWS GovCloud (US-West) – 5 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual Amazon Rekognition Image data plane operations:<br>• CompareFaces<br>• DetectFaces | • US East (Ohio) Region – 25<br>• US East (N. Virginia) Region – 100<br>• US West (N. California) Region – 25<br>• US West (Oregon) Region – 100<br>• Asia Pacific (Mumbai) Region – 25<br>• Asia Pacific (Seoul) Region – 25<br>• Asia Pacific (Singapore) Region – 25<br>• Asia Pacific (Sydney) Region – 25<br>• Asia Pacific (Tokyo) Region – 25<br>• |

| Resource | Default |
| --- | --- |
| | Canada (Central) – 25 (For supported operations, see Canada (Central) Region).<br>• Europe (Frankfurt) Region – 25<br>• Europe (Spain) Region – 25<br>• South America (São Paulo) Region – 25<br>• Asia Pacific (Malaysia) Region – 25<br>• Asia Pacific (Thailand) Region – 25<br>• Europe (Ireland) Region – 100<br>• Europe (London) Region – 25<br>• Israel (Tel Aviv) Region – 25 (For supported operations, see Israel (Tel Aviv) Region) |

| Resource | Default |
| --- | --- |
| | • AWS GovCloud (US-West) – 25 |
| Transactions per second per account for the Amazon Rekogniti on Image data plane operation:<br>• AssociateFaces<br>• DisassociateFaces | In each Region that Amazon Rekognition Image supports – 5 |
| Transactions per second per account for the personal protectiv e equipment data plane operation:<br>• DetectProtectiveEquipment | In each Region that Amazon Rekognition Image supports – 5 |
| Transactions per second per account for individual Amazon Rekognition Image control plane operations:<br>• CreateCollection<br>• DeleteCollection<br>• DeleteFaces<br>• DescribeCollection<br>• ListCollections<br>• CreateUser<br>• DeleteUser<br>• ListUsers | In each Region that Amazon Rekognition Image supports – 5 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual stored video start operations:<br>• StartCelebrityRecognition<br>• StartContentModeration<br>• StartFaceDetection<br>• StartFaceSearch<br>• StartLabelDetection<br>• StartPersonTracking<br>• StartTextDetection<br>• StartSegmentDetection | In each Region that Amazon Rekognition Video supports – 5<br>`StartCelebrityRecognition` is not available in AWS GovCloud (US). |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual Amazon Rekognition Video stored video get operations:<br>• GetCelebrityRecognition<br>• GetContentModeration<br>• GetFaceDetection<br>• GetFaceSearch<br>• GetLabelDetection<br>• GetPersonTracking<br>• GetTextDetection<br>• GetSegmentDetection | • US East (Ohio) Region – 5<br>• US East (N. Virginia) Region – 20<br>• US West (N. California) Region – 5<br>• US West (Oregon) Region – 20<br>• Asia Pacific (Mumbai) Region – 5<br>• Asia Pacific (Seoul) Region – 5<br>• Asia Pacific (Singapore) Region – 5<br>• Asia Pacific (Sydney) Region – 5<br>• Asia Pacific (Tokyo) Region – 5<br>• |

| Resource | Default |
| --- | --- |
| | Europe (Frankfurt) Region – 5<br>• South America (São Paulo) Region – 5<br>• Asia Pacific (Malaysia) Region – 5<br>• Asia Pacific (Thailand) Region – 5<br>• Europe (Ireland) Region – 20<br>• Europe (London) Region – 5<br>• AWS GovCloud (US-West) – 20 (`GetCelebrityRecognition` is not available in this Region.) |
| Maximum number of concurrent stored video jobs per account | 20 |
| Transactions per second per account for individual bulk analysis start operations: StartMediaAnalysisJob | In each Region that Amazon Rekognition bulk analysis supports: 5 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual bulk analysis list operations: ListMediaAnalysisJob | In each Region that Amazon Rekognition bulk analysis supports: 5 |
| Transactions per second per account for individual bulk analysis get operations: GetMediaAnalysisJob | In each Region that Amazon Rekognition bulk analysis supports: 20 |

| Resource | Default |
| --- | --- |
| Maximum number of concurrent bulk analysis jobs per account, with regard to region: | • US East (N. Virginia) Region – 20<br>• US East (Ohio) Region – 5<br>• US West (Oregon) Region – 20<br>• Europe (Ireland) Region – 20<br>• Asia Pacific (Mumbai) Region – 5<br>• Asia Pacific (Seoul) Region – 5<br>• Asia Pacific (Singapore) Region – 5<br>• Asia Pacific (Sydney) Region – 5<br>• Asia Pacific (Tokyo) Region – 5<br>• |

| Resource | Default |
| --- | --- |
| | Europe (Frankfurt) Region – 5<br>• Europe (London) Region – 5 |
| Maximum number of streaming video stream processors per account that can simultaneously exist | In each Region that Amazon Rekognition Video supports – 10,000 |
| Maximum number of face search stream processors per account that can be processed concurrently | In each Region that Amazon Rekognition Video supports face search stream processors – 10 |
| Maximum number of label detection stream processors per account that can be processed concurrently | • US East (N. Virginia) – 200<br>• US East (Ohio) – 40<br>• US West (Oregon) – 200<br>• Asia Pacific (Mumbai) – 40<br>• Europe (Ireland) – 40 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual streaming video operations:<br>• CreateStreamProcessor<br>• DeleteStreamProcessor<br>• DescribeStreamProcessor<br>• StartStreamProcessor<br>• UpdateStreamProcessor | In each Region that Amazon Rekognition Video supports – 20 |
| Transactions per second per account for stop streaming video operations:<br>• StopStreamProcessor | In each Region that Amazon Rekognition Video supports – 1 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for Amazon Rekognition Face Liveness API operations:<br>• CreateFaceLivenessSession<br>• GetFaceLivenessSessionResults<br>• StartFaceLivenessSession | • US East (N. Virginia) – 25<br>• Europe (Ireland) – 5<br>• US West (Oregon) – 25<br>• Asia Pacific (Mumbai) – 5<br>• Asia Pacific (Tokyo) Region – 5<br>• South America (São Paulo) Region – 5<br>• Asia Pacific (Malaysia) Region – 5<br>• Asia Pacific (Thailand) Region – 5 |

| Resource | Default |
| --- | --- |
| Number of concurrent Amazon Rekognition Face Liveness sessions per account, created with StartFaceLivenessSession.<br><br>To determine your required concurrent sessions quota, multiply the estimated session length by your estimated TPS. (Ex. 10 seconds x 5 TPS = 50). | • US East (N. Virginia) – 75<br>• Europe (Ireland) – 15<br>• US West (Oregon) – 75<br>• Asia Pacific (Mumbai) – 15<br>• Asia Pacific (Tokyo) Region – 15<br>• South America (São Paulo) Region – 15<br>• Asia Pacific (Malaysia) Region – 15<br>• Asia Pacific (Thailand) Region – 15 |
| Transactions per second per account for list streaming video operations:<br><br>• ListStreamProcessors | In each Region that Amazon Rekognition Video supports – 5 |
| Transactions per second per account for resource tagging operations:<br><br>• ListTagsForResource<br>• TagResource<br>• UntagResource | In each Region that Amazon Rekognition Image supports – 10 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual Amazon Rekognition  Custom Label data plane operations:<br><br>• DetectCustomLabels | In all Regions that Amazon Rekognition Custom Labels supports – 50 |

| Resource | Default |
| --- | --- |
| Transactions per second per account for individual Amazon Rekognition  Custom Labels control plane operations:<br>• CopyProjectVersion<br>• CreateDataset<br>• CreateProject<br>• CreateProjectVersion<br>• DeleteDataset<br>• DeleteProject<br>• DeleteProjectPolicy<br>• DeleteProjectVersion<br>• DescribeDataset<br>• DescribeProjects<br>• DescribeProjectVersions<br>• DistributeDatasetEntries<br>• ListDatasetEntries<br>• ListDatasetLabels<br>• ListProjectPolicies<br>• PutProjectPolicy<br>• StartProjectVersion | In each Region that Amazon Rekognition Custom Labels supports – 5 |

| Resource | Default |
| --- | --- |
| • StopProjectVersion<br>• UpdateDatasetEntries | |
| Maximum number of Amazon Rekognition Custom Labels projects per account. | 100 |
| Maximum number of Amazon Rekognition Custom Labels models per project. | 100 |
| Maximum number of concurrent Amazon Rekognition Custom Labels training jobs per account. | • All Regions except Asia Pacific (Sydney) – 2<br>• Asia Pacific (Sydney) – 1 |
| Maximum number of concurrently running Amazon Rekogniti on Custom Labels models per account. | 2 |
| Maximum inference units per started model. | 5 |
| Maximum number of images per dataset. | 250,000 |

For more information, see Guidelines and quotas in Amazon Rekognition in the *Amazon Rekognition Developer Guide*.

## Amazon Relational Database Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Amazon RDS

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | rds.us-east-2.amazonaws.com | HTTPS | |
| | | rds-fips.us-east-2.api.aws | HTTPS | |
| | | rds.us-east-2.api.aws | HTTPS | |
| | | rds-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | rds.us-east-1.amazonaws.com | HTTPS | |
| | | rds-fips.us-east-1.api.aws | HTTPS | |
| | | rds-fips.us-east-1.amazonaws.com | HTTPS | |
| | | rds.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | rds.us-west-1.amazonaws.com | HTTPS | |
| | | rds.us-west-1.api.aws | HTTPS | |
| | | rds-fips.us-west-1.amazonaws.com | HTTPS | |
| | | rds-fips.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | rds.us-west-2.amazonaws.com | HTTPS | |
| | | rds-fips.us-west-2.amazonaws.com | HTTPS | |
| | | rds.us-west-2.api.aws | HTTPS | |
| | | rds-fips.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | rds.af-south-1.amazonaws.com<br>rds.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | rds.ap-east-1.amazonaws.com<br>rds.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | rds.ap-south-2.amazonaws.com<br>rds.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | rds.ap-southeast-3.amazonaws.com<br>rds.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | rds.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | rds.ap-southeast-4.amazonaws.com<br>rds.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | rds.ap-south-1.amazonaws.com<br>rds.ap-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | rds.ap-southeast-6.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | rds.ap-northeast-3.amazonaws.com<br>rds.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | rds.ap-northeast-2.amazonaws.com<br>rds.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | rds.ap-southeast-1.amazonaws.com<br>rds.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | rds.ap-southeast-2.amazonaws.com<br>rds.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | rds.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | rds.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | rds.ap-northeast-1.amazonaws.com<br>rds.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | rds.ca-central-1.amazonaws.com | HTTPS | |
| | | rds.ca-central-1.api.aws | HTTPS | |
| | | rds-fips.ca-central-1.api.aws | HTTPS | |
| | | rds-fips.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | rds.ca-west-1.amazonaws.com | HTTPS | |
| | | rds-fips.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | rds.eu-central-1.amazonaws.com | HTTPS | |
| | | rds.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | rds.eu-west-1.amazonaws.com | HTTPS | |
| | | rds.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | rds.eu-west-2.amazonaws.com | HTTPS | |
| | | rds.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | rds.eu-south-1.amazonaws.com | HTTPS | |
| | | rds.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | rds.eu-west-3.amazonaws.com | HTTPS | |
| | | rds.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | rds.eu-south-2.amazonaws.com | HTTPS | |
| | | rds.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | rds.eu-north-1.amazonaws.com<br>rds.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | rds.eu-central-2.amazonaws.com<br>rds.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | rds.il-central-1.amazonaws.com<br>rds.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | rds.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | rds.me-south-1.amazonaws.com<br>rds.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | rds.me-central-1.amazonaws.com<br>rds.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | rds.sa-east-1.amazonaws.com<br>rds.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | rds.us-gov-east-1.amazonaws.com<br>rds.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | rds.us-gov-west-1.amazonaws.com<br>rds.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

## Amazon RDS Performance Insights

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | pi.us-east-2.amazonaws.com | HTTPS | |
| | | pi-fips.us-east-2.api.aws | HTTPS | |
| | | pi-fips.us-east-2.amazonaws.com | HTTPS | |
| | | pi.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | pi.us-east-1.amazonaws.com | HTTPS | |
| | | pi-fips.us-east-1.api.aws | HTTPS | |
| | | pi-fips.us-east-1.amazonaws.com | HTTPS | |
| | | pi.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | pi.us-west-1.amazonaws.com | HTTPS | |
| | | pi-fips.us-west-1.api.aws | HTTPS | |
| | | pi-fips.us-west-1.amazonaws.com | HTTPS | |
| | | pi.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | pi.us-west-2.amazonaws.com | HTTPS | |
| | | pi-fips.us-west-2.api.aws | HTTPS | |
| | | pi-fips.us-west-2.amazonaws.com | HTTPS | |
| | | pi.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | pi.af-south-1.amazonaws.com | HTTPS | |
| | | pi.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | pi.ap-east-1.amazonaws.com<br>pi.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | pi.ap-south-2.amazonaws.com<br>pi.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | pi.ap-southeast-3.amazonaws.com<br>pi.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | pi.ap-southeast-5.amazonaws.com<br>pi.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | pi.ap-southeast-4.amazonaws.com<br>pi.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | pi.ap-south-1.amazonaws.com<br>pi.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | pi.ap-southeast-6.amazonaws.com<br>pi.ap-southeast-6.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | pi.ap-northeast-3.amazonaws.com<br>pi.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | pi.ap-northeast-2.amazonaws.com<br>pi.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | pi.ap-southeast-1.amazonaws.com<br>pi.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | pi.ap-southeast-2.amazonaws.com<br>pi.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | pi.ap-east-2.amazonaws.com<br>pi.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | pi.ap-southeast-7.amazonaws.com<br>pi.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | pi.ap-northeast-1.amazonaws.com<br>pi.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | pi.ca-central-1.amazonaws.com | HTTPS |
| | | pi-fips.ca-central-1.api.aws | HTTPS |
| | | pi-fips.ca-central-1.amazonaws.com | HTTPS |
| | | pi.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | pi.ca-west-1.amazonaws.com | HTTPS |
| | | pi-fips.ca-west-1.api.aws | HTTPS |
| | | pi-fips.ca-west-1.amazonaws.com | HTTPS |
| | | pi.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | pi.eu-central-1.amazonaws.com | HTTPS |
| | | pi.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | pi.eu-west-1.amazonaws.com | HTTPS |
| | | pi.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | pi.eu-west-2.amazonaws.com | HTTPS |
| | | pi.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | pi.eu-south-1.amazonaws.com | HTTPS |
| | | pi.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | pi.eu-west-3.amazonaws.com | HTTPS |
| | | pi.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | pi.eu-south-2.amazonaws.com | HTTPS |
| | | pi.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | pi.eu-north-1.amazonaws.com | HTTPS |
| | | pi.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | pi.eu-central-2.amazonaws.com | HTTPS |
| | | pi.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | pi.il-central-1.amazonaws.com | HTTPS |
| | | pi.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | pi.mx-central-1.amazonaws.com | HTTPS |
| | | pi.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | pi.me-south-1.amazonaws.com | HTTPS |
| | | pi.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | pi.me-central-1.amazonaws.com | HTTPS |
| | | pi.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | pi.sa-east-1.amazonaws.com | HTTPS |
| | | pi.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | pi.us-gov-east-1.amazonaws.com | HTTPS |
| | | pi-fips.us-gov-east-1.api.aws | HTTPS |
| | | pi-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | pi.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | pi.us-gov-west-1.amazonaws.com | HTTPS | |
| | | pi-fips.us-gov-west-1.api.aws | HTTPS | |
| | | pi-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | pi.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Authorizations per DB security group | Each supported Region: 20 | No | Number of security group authorizations per DB security group |
| Custom endpoints per DB cluster | Each supported Region: 5 | Yes | The maximum number of custom endpoints that you can create per Aurora DB cluster in this account in the current Region. This value reflects the highest number of custom endpoints in a DB cluster in the account. Other DB clusters in the account might have a lower number of custom endpoints. |
| Custom engine versions | Each supported Region: 40 | Yes | The maximum number of custom engine versions |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | allowed in this account in the current Region |
| DB cluster parameter groups | Each supported Region: 50 | Yes | The maximum number of DB cluster parameter groups |
| DB clusters | Each supported Region: 40 | Yes | The maximum number of Aurora clusters allowed in this account in the current Region |
| DB instances | ap-south-1: 20<br><br>Each of the other supported Regions: 40 | Yes | The maximum number of DB instances allowed in this account in the current Region |
| DB shard groups | Each supported Region: 5 | Yes | The maximum number of DB shard groups for Aurora Limitless Database in this account in the current Region |
| DB subnet groups | ap-south-1: 20<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of DB subnet groups |
| Data API HTTP request body size | Each supported Region: 4 Megabytes | No | The maximum size allowed for the HTTP request body. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Data API maximum concurrent cluster-secret pairs | Each supported Region: 30 | No | The maximum number of unique pairs of Aurora Serverless v1 DB clusters and secrets in concurrent Data API requests for this account in the current AWS Region. |
| Data API maximum concurrent requests | Each supported Region: 500 | No | The maximum number of Data API requests to an Aurora Serverless v1 DB cluster that use the same secret and can be processed at the same time. Additional requests are queued and processed as in-process requests complete. |
| Data API maximum result set size | Each supported Region: 1 Megabytes | No | The maximum size of the database result set that can be returned by the Data API. |
| Data API maximum size of JSON response string | Each supported Region: 10 Megabytes | No | The maximum size of the simplified JSON response string returned by the RDS Data API. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Data API requests per second | Each supported Region: 1,000 per second | No | The maximum number of requests to the Data API per second allowed for this account in the current AWS Region. This quota only applies to Amazon Aurora Serverless v1 clusters. |
| Event subscriptions | Each supported Region: 20 | Yes | The maximum number of event subscriptions |
| IAM roles per DB cluster | Each supported Region: 5 | Yes | The maximum number of IAM roles associated with a DB cluster |
| IAM roles per DB instance | Each supported Region: 5 | Yes | The maximum number of IAM roles associated with a DB instance |
| Integrations | Each supported Region: 100 | No | The maximum number of integrations allowed in this account in the current AWS Region |
| Manual DB cluster snapshots | Each supported Region: 100 | Yes | The maximum number of manual DB cluster snapshots |
| Manual DB instance snapshots | Each supported Region: 100 | Yes | The maximum number of manual DB instance snapshots |
| Option groups | Each supported Region: 20 | Yes | The maximum number of option groups |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Parameter groups | ap-south-1: 20<br><br>Each of the other supported Regions: 50 | Yes | The maximum number of parameter groups |
| Proxies | Each supported Region: 20 | Yes | The maximum number of proxies allowed in this account in the current AWS Region |
| Read replicas per primary | Each supported Region: 15 | Yes | The maximum number of read replicas per primary DB instance. This quota cant be adjusted for Amazon Aurora. |
| Reserved DB instances | ap-south-1: 20<br><br>Each of the other supported Regions: 40 | Yes | The maximum number of reserved DB instances allowed in this account in the current AWS Region |
| Security groups | ap-south-1: 20<br><br>Each of the other supported Regions: 25 | Yes | The maximum number of DB security groups |
| Subnets per DB subnet group | Each supported Region: 20 | No | The maximum number of subnets per DB subnet group |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Total storage for all DB instances | Each supported Region: 100,000 | Yes | The maximum total storage (in GB) on EBS volumes for all Amazon RDS DB instances added together. This quota does not apply to Amazon Aurora, which has a maximum cluster volume of 128 TiB for each DB cluster. |

# AWS re:Post Private endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | repostspace.us-east-1.amazonaws.com repostspace.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | repostspace.us-west-2.amazonaws.com repostspace.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | repostspace.ap-southeast-1.amazonaws.com repostspace.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | repostspace.ap-southeast-2.amazonaws.com repostspace.ap-southeast-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | repostspace.ca-central-1.amazonaws.com repostspace.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | repostspace.eu-central-1.amazonaws.com repostspace.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | repostspace.eu-west-1.amazonaws.com repostspace.eu-west-1.api.aws | HTTPS |

## Service quotas

The AWS re:Post Private service has the following quotas:

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Free private re:Post size | Each supported Region: 10 Gigabytes | No | The maximum size (in GB) of a free private re:Post. |
| Number of private re:Posts | Each supported Region: 3 | No | The maximum number of private re:Posts in this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateSpace API requests | Each supported Region: 1 | No | The maximum number of CreateSpace calls that you can make per second. |
| Rate of DeleteSpace API requests | Each supported Region: 1 | No | The maximum number of DeleteSpace calls that you can make per second. |
| Rate of DeregisterAdmin API requests | Each supported Region: 100 | No | The maximum number of DeregisterAdmin calls that you can make per second. |
| Rate of GetSpace API requests | Each supported Region: 10 | No | The maximum number of GetSpace calls that you can make per second. |
| Rate of ListSpaces API requests | Each supported Region: 10 | No | The maximum number of ListSpaces calls that you can make per second. |
| Rate of ListTagsForResource API requests | Each supported Region: 10 | No | The maximum number of ListTagsForResource calls that you can make per second. |
| Rate of RegisterAdmin API requests | Each supported Region: 100 | No | The maximum number of RegisterAdmin calls that you can make per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of SendInvites API requests | Each supported Region: 1 | No | The maximum number of SendInvites calls that you can make per second. |
| Rate of TagResource API requests | Each supported Region: 10 | No | The maximum number of TagResource calls that you can make per second. |
| Rate of UntagResource API requests | Each supported Region: 10 | No | The maximum number of UntagResource calls that you can make per second. |
| Rate of UpdateSpace API requests | Each supported Region: 10 | No | The maximum number of UpdateSpace calls that you can make per second. |
| Standard private re:Post size | Each supported Region: 100 Gigabytes | No | The maximum size (in GB) of a standard private re:Post. |
| Token refill rate for CreateSpace API requests | Each supported Region: 1 | No | CreateSpace throttle token refill rate per second. |
| Token refill rate for DeleteSpace API requests | Each supported Region: 1 | No | DeleteSpace throttle token refill rate per second. |
| Token refill rate for DeregisterAdmin API requests | Each supported Region: 10 | No | DeregisterAdmin throttle token refill rate per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Token refill rate for GetSpace API requests | Each supported Region: 10 | No | GetSpace throttle token refill rate per second. |
| Token refill rate for ListSpaces API requests | Each supported Region: 10 | No | ListSpaces throttle token refill rate per second. |
| Token refill rate for ListTagsForResource API requests | Each supported Region: 10 | No | ListTagsForResource throttle token refill rate per second. |
| Token refill rate for RegisterAdmin API requests | Each supported Region: 10 | No | RegisterAdmin throttle token refill rate per second. |
| Token refill rate for SendInvites API requests | Each supported Region: 1 | No | SendInvites throttle token refill rate per second. |
| Token refill rate for TagResource API requests | Each supported Region: 10 | No | TagResource throttle token refill rate per second. |
| Token refill rate for UntagResource API requests | Each supported Region: 10 | No | UntagResource throttle token refill rate per second. |
| Token refill rate for UpdateSpace API requests | Each supported Region: 10 | No | UpdateSpace throttle token refill rate per second. |

For more information, see re:Post Private quotas.

# AWS Resilience Hub endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | resiliencehub.us-east-2.amazonaws.com | HTTPS |
| | | resiliencehub.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | resiliencehub.us-east-1.amazonaws.com | HTTPS |
| | | resiliencehub.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | resiliencehub.us-west-1.amazonaws.com | HTTPS |
| | | resiliencehub.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | resiliencehub.us-west-2.amazonaws.com | HTTPS |
| | | resiliencehub.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | resiliencehub.af-south-1.amazonaws.com | HTTPS |
| | | resiliencehub.af-south-1.api.aws | HTTPS |
| Asia Pacific | ap-east-1 | resiliencehub.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| (Hong Kong) | | resiliencehub.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | resiliencehub.ap-south-1.amazonaws.com<br>resiliencehub.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | resiliencehub.ap-northeast-2.amazonaws.com<br>resiliencehub.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | resiliencehub.ap-southeast-1.amazonaws.com<br>resiliencehub.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | resiliencehub.ap-southeast-2.amazonaws.com<br>resiliencehub.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | resiliencehub.ap-northeast-1.amazonaws.com<br>resiliencehub.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | resiliencehub.ca-central-1.amazonaws.com<br>resiliencehub.ca-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | resiliencehub.eu-central-1.amazonaws.com<br>resiliencehub.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | resiliencehub.eu-west-1.amazonaws.com<br>resiliencehub.eu-west-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | resiliencehub.eu-west-2.amazonaws.com | HTTPS |
| | | resiliencehub.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | resiliencehub.eu-south-1.amazonaws.com | HTTPS |
| | | resiliencehub.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | resiliencehub.eu-west-3.amazonaws.com | HTTPS |
| | | resiliencehub.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | resiliencehub.eu-north-1.amazonaws.com | HTTPS |
| | | resiliencehub.eu-north-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | resiliencehub.me-south-1.amazonaws.com | HTTPS |
| | | resiliencehub.me-south-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | resiliencehub.sa-east-1.amazonaws.com | HTTPS |
| | | resiliencehub.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | resiliencehub.us-gov-east-1.amazonaws.com | HTTPS |
| | | resiliencehub-fips.us-gov-east-1.api.aws | HTTPS |
| | | resiliencehub-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | resiliencehub.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | resiliencehub.us-gov-west-1.amazonaws.com<br>resiliencehub-fips.us-gov-west-1.api.aws<br>resiliencehub-fips.us-gov-west-1.amazonaws.com<br>resiliencehub.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of AWS CloudFormation stacks to import | Each supported Region: 20 | Yes | The maximum number of AWS CloudFormation stacks an AWS account can import for a given application |
| Number of Application Components per application | Each supported Region: 200 | Yes | The maximum number of Application Components for a given application |
| Number of Application Components per resource | Each supported Region: 5 | Yes | The maximum number of Application Components an AWS account can have for a given resource |
| Number of Compliance Readiness Policies | Each supported Region: 10 | Yes | The maximum number of compliance readiness policies an AWS account can create in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of EKS clusters to import | Each supported Region: 5 | Yes | The maximum number of EKS clusters an AWS account can import for a given application. |
| Number of ResilienceHubV2 assumptions per service | Each supported Region: 20 | No | The maximum number of ResilienceHubV2 assumptions per service |
| Number of ResilienceHubV2 concurrent assessments per service | Each supported Region: 1 | No | The maximum number of ResilienceHubV2 concurrent assessments per service |
| Number of ResilienceHubV2 cross-account role ARNs per service | Each supported Region: 5 | No | The maximum number of ResilienceHubV2 cross-account role ARNs per service |
| Number of ResilienceHubV2 input sources per service | Each supported Region: 20 | Yes | The maximum number of ResilienceHubV2 input sources per service |
| Number of ResilienceHubV2 policies per account | Each supported Region: 100 | No | The maximum number of ResilienceHubV2 policies an account can create in the current Region |
| Number of ResilienceHubV2 regions per service | Each supported Region: 5 | No | The maximum number of ResilienceHubV2 regions per service |
| Number of ResilienceHubV2 resources per service | Each supported Region: 500 | Yes | The maximum number of ResilienceHubV2 resources per service |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of ResilienceHubV2 service functions per service | Each supported Region: 20 | No | The maximum number of ResilienceHubV2 service functions per service |
| Number of ResilienceHubV2 services per account | Each supported Region: 100 | Yes | The maximum number of ResilienceHubV2 services an account can create in the current Region |
| Number of ResilienceHubV2 services per user journey | Each supported Region: 20 | No | The maximum number of ResilienceHubV2 services per user journey |
| Number of ResilienceHubV2 systems per account | Each supported Region: 500 | Yes | The maximum number of ResilienceHubV2 systems an account can create in the current Region |
| Number of ResilienceHubV2 systems per service | Each supported Region: 20 | No | The maximum number of ResilienceHubV2 systems associated per service |
| Number of ResilienceHubV2 tags per input source | Each supported Region: 10 | No | The maximum number of ResilienceHubV2 tags per input source |
| Number of ResilienceHubV2 user journeys per system | Each supported Region: 20 | No | The maximum number of ResilienceHubV2 user journeys per system |
| Number of Resiliency Policies | Each supported Region: 10 | Yes | The maximum number of resiliency policies an AWS account can create in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of Terraform state files to import | Each supported Region: 20 | Yes | The maximum number of Terraform state files an AWS account can import for a given application |
| Number of applications | Each supported Region: 50 | Yes | The maximum number of applications an AWS account can create in the current Region |
| Number of assessments per application per month | Each supported Region: 200 | Yes | The maximum number of assessments an AWS account can run for a given application in a given month |
| Number of concurrent assessments per account | Each supported Region: 20 | No | The maximum number of concurrent assessments an AWS account can run |
| Number of concurrent assessments per application | Each supported Region: 1 | No | The maximum number of concurrent assessments an AWS account can run for a given application |
| Number of concurrent recommendation templates per account | Each supported Region: 100 | No | The maximum number of recommendation templates an AWS account can concurrently create |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of concurrent recommendation templates per application | Each supported Region: 1 | No | The maximum number of recommendation templates an AWS account can create for a given application |
| Number of metric exports per month | Each supported Region: 300 | Yes | The maximum number of metric exports an AWS account can run in a given month |
| Number of namespaces to import for an EKS cluster | Each supported Region: 10 | Yes | The maximum number of namespaces per EKS cluster an AWS account can import. |
| Number of recommendation templates per application per month | Each supported Region: 100 | Yes | The maximum number of recommendation templates an AWS account can create for a given application in a given month |
| Number of resources per application | Each supported Region: 400 | Yes | The maximum number of resources for a given application |
| ResilienceHubV2 Terraform template maximum size | Each supported Region: 26,214,400 | No | The maximum size of a ResilienceHubV2 Terraform template in bytes |
| ResilienceHubV2 default pagination size | Each supported Region: 100 | No | The ResilienceHubV2 default pagination size |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ResilienceHubV2 event log retention period in days | Each supported Region: 730 | No | The ResilienceHubV2 event log retention period in days |
| Retention period of past assessments/recommendations in days | Each supported Region: 365 | No | Retention period of past assessments/recommendations in days |
| Retention period of past recommendation templates in days | Each supported Region: 365 | No | Retention period of past recommendation templates in days |
| Template size in bytes | Each supported Region: 51,200 | Yes | The maximum size of a template in bytes |
| Terraform state file maximum size | Each supported Region: 4,194,305 | Yes | The maximum import size limit for terraform state files |

## AWS Resource Explorer

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | resource-explorer-2.us-east-2.amazonaws.com<br>resource-explorer-2-fips.us-east-2.amazonaws.com<br>resource-explorer-2-fips.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | resource-explorer-2.us-east-1.amazonaws.com<br>resource-explorer-2-fips.us-east-1.amazonaws.com<br>resource-explorer-2-fips.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | resource-explorer-2.us-west-1.amazonaws.com<br>resource-explorer-2-fips.us-west-1.amazonaws.com<br>resource-explorer-2-fips.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | resource-explorer-2.us-west-2.amazonaws.com<br>resource-explorer-2-fips.us-west-2.amazonaws.com<br>resource-explorer-2-fips.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | resource-explorer-2.af-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | resource-explorer-2.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | resource-explorer-2.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | resource-explorer-2.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | resource-explorer-2.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | resource-explorer-2.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | resource-explorer-2.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | resource-explorer-2.ap-southeast-6.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | resource-explorer-2.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | resource-explorer-2.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | resource-explorer-2.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | resource-explorer-2.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | resource-explorer-2.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | resource-explorer-2.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | resource-explorer-2.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | resource-explorer-2.ca-central-1.amazonaws.com<br>resource-explorer-2-fips.ca-central-1.amazonaws.com<br>resource-explorer-2-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | resource-explorer-2.ca-west-1.amazonaws.com<br>resource-explorer-2-fips.ca-west-1.amazonaws.com<br>resource-explorer-2-fips.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | resource-explorer-2.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | resource-explorer-2.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | resource-explorer-2.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | resource-explorer-2.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | resource-explorer-2.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | resource-explorer-2.eu-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | resource-explorer-2.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | resource-explorer-2.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | resource-explorer-2.il-central-1.amazonaws.com | HTTPS |
| Mexico (Central) | mx-central-1 | resource-explorer-2.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | resource-explorer-2.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | resource-explorer-2.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | resource-explorer-2.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | resource-explorer-2.us-gov-east-1.amazonaws.com<br>resource-explorer-2-fips.us-gov-east-1.amazonaws.com<br>resource-explorer-2-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | resource-explorer-2.us-gov-west-1.amazonaws.com<br><br>resource-explorer-2-fips.us-gov-west-1.amazonaws.com<br><br>resource-explorer-2-fips.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS<br><br>HTTPS | |

## Service quotas

Your AWS account has default quotas for each AWS service. Unless otherwise noted, quotas are Region-specific. You can request increases for some quotas, and other quotas cannot be increased.

To view the quotas for AWS Resource Explorer, open the Service Quotas console. In the navigation pane, choose **AWS services** and select **AWS Resource Explorer**.

To request a quota increase, see Requesting a Quota Increase in the *Service Quotas User Guide*. If the quota is not yet available in Service Quotas, use the limit increase form.

The following quotas are the defaults for Resource Explorer.

| Maximum value quotas | Default value |
|---|---|
| Number of views in an AWS Region | 10 |

| Rate limits for operations | Default value |
|---|---|
| Maximum Search operations per second | 5 |
| Maximum non-Search operations per second | 3 |
| Maximum Search operations in aggregator Region per month | 10,000 |

| Rate limits for operations | Default value |
| --- | --- |
| Maximum Search operations in local Regions per month | 500 |

# AWS Resource Groups and Tagging endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## AWS Resource Groups

### Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | resource-groups.us-east-2.amazonaws.com | HTTPS |
| | | resource-groups-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | resource-groups.us-east-1.amazonaws.com | HTTPS |
| | | resource-groups-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | resource-groups.us-west-1.amazonaws.com | HTTPS |
| | | resource-groups-fips.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | resource-groups.us-west-2.amazonaws.com<br><br>resource-groups-fips.us-west-2.amazonaws.com | HTTPS<br><br>HTTPS | |
| Africa (Cape Town) | af-south-1 | resource-groups.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | resource-groups.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | resource-groups.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | resource-groups.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | resource-groups.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | resource-groups.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | resource-groups.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | resource-groups.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | resource-groups.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | resource-groups.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | resource-groups.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | resource-groups.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | resource-groups.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | resource-groups.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | resource-groups.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | resource-groups.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | resource-groups.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | resource-groups.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | resource-groups.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | resource-groups.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | resource-groups.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | resource-groups.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | resource-groups.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | resource-groups.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | resource-groups.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | resource-groups.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | resource-groups.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | resource-groups.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | resource-groups.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | resource-groups.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | resource-groups.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | resource-groups.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas for Resource Groups

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Resource groups per account | Each supported Region: 100 | Yes | The maximum number of resource groups that you can create in this account. A resource |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| | | | group is a collection of AWS resources that match a specific criteria. |

# Tagging

## Service endpoints for AWS Resource Groups Tagging API

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | tagging.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | tagging.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | tagging.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | tagging.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | tagging.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | tagging.ap-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | tagging.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | tagging.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | tagging.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | tagging.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | tagging.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | tagging.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | tagging.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | tagging.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | tagging.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | tagging.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | tagging.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | tagging.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | tagging.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | tagging.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | tagging.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | tagging.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | tagging.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | tagging.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | tagging.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | tagging.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | tagging.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | tagging.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | tagging.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | tagging.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | tagging.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | tagging.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | tagging.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | tagging.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | tagging.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | tagging.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas for Tagging (Tag Editor and Resource Groups Tagging API)

| Name | Default | |
|---|---|---|
| Tags attached per resource | 50 user created tags (AWS created tags don't count against this limit) | |
| Tag key name | Minimum of 1, maximum 128 Unicode characters in UTF-8.<br><br>Allowed characters include: letters, numbers, spaces, and the following characters:<br><br>`_ . : / = + - @`<br><br>Key names can't begin with `aws:` because that is reserved. | |

| Name | Default |
| --- | --- |
| | **Note**<br>Some AWS services have some additional character or length restrictions. For details, see the documentation for the specific service. |
| Tag values | Minimum of 0, maximum of 256 Unicode characters in UTF-8.<br><br>Allowed characters include: letters, numbers, spaces, and the following characters:<br><br>`_ . : / = + - @`<br><br>**Note**<br>Some AWS services have some additional character or length restrictions. For details, see the documentation for the specific service. |
| Rate of calling the GetResources API operation | Maximum of 15 calls per second |

| Name | Default | |
|---|---|---|
| Rate of calling the following API operations:<br>• TagResources<br>• UntagResources<br>• GetTagKeys<br>• GetTagValues | Maximum of 5 calls per second | |

**Note**

These limits are currently not adjustable using the Service Quotas console. Contact AWS Support.

# Red Hat OpenShift Service on AWS endpoints and quotas

Red Hat OpenShift Service on AWS (ROSA) currently does not support programmatic access to service endpoints via the AWS CLI. The following are the AWS Regions and service quotas for this service. Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

## Service endpoints

| Region name | Region | Available for ROSA with hosted control planes | Available for ROSA classic |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | Yes | Yes |
| US East (Ohio) | us-east-2 | Yes | Yes |
| US West (N. California) | us-west-1 | No | Yes |

| Region name | Region | Available for ROSA with hosted control planes | Available for ROSA classic |
|---|---|---|---|
| US West (Oregon) | us-west-2 | Yes | Yes |
| Africa (Cape Town) | af-south-1 | Yes | Yes |
| Asia Pacific (Hong Kong) | ap-east-1 | Yes | Yes |
| Asia Pacific (Hyderabad) | ap-south-2 | Yes | Yes |
| Asia Pacific (Jakarta) | ap-southeast-3 | Yes | Yes |
| Asia Pacific (Melbourne) | ap-southeast-4 | Yes | Yes |
| Asia Pacific (Mumbai) | ap-south-1 | Yes | Yes |
| Asia Pacific (Osaka) | ap-northeast-3 | Yes | Yes |
| Asia Pacific (Seoul) | ap-northeast-2 | Yes | Yes |
| Asia Pacific (Singapore) | ap-southeast-1 | Yes | Yes |
| Asia Pacific (Sydney) | ap-southeast-2 | Yes | Yes |
| Asia Pacific (Tokyo) | ap-northeast-1 | Yes | Yes |
| Canada (Central) | ca-central-1 | Yes | Yes |
| Europe (Frankfurt) | eu-central-1 | Yes | Yes |
| Europe (Ireland) | eu-west-1 | Yes | Yes |
| Europe (London) | eu-west-2 | Yes | Yes |
| Europe (Milan) | eu-south-1 | Yes | Yes |

| Region name | Region | Available for ROSA with hosted control planes | Available for ROSA classic |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | Yes | Yes |
| Europe (Spain) | eu-south-2 | Yes | Yes |
| Europe (Stockholm) | eu-north-1 | Yes | Yes |
| Europe (Zurich) | eu-central-2 | Yes | Yes |
| Middle East (Bahrain) | me-south-1 | Yes | Yes |
| Middle East (UAE) | me-central-1 | Yes | Yes |
| South America (São Paulo) | sa-east-1 | Yes | Yes |
| AWS GovCloud (US-East) | us-gov-east-1 | No | Yes |
| AWS GovCloud (US-West) | us-gov-west-1 | No | Yes |

## Service quotas

ROSA uses service quotas for Amazon EC2, Amazon Virtual Private Cloud, Amazon Elastic Block Store, and Elastic Load Balancing. These quotas are listed in the corresponding namespace in the Service Quotas console. Although most default values are suitable for most workloads, you may need to request a quota increase for the following cases.

- ROSA classic clusters require a minimum of 100 vCPUs for cluster creation, availability, and upgrades. The default value for Running On-Demand Standard (A, C, D, H, I, M, R, T, Z) Amazon EC2 instances is 5 vCPUs. If you do not increase this quota, cluster creation fails.
- Some optional cluster configuration features, such as custom security groups, may require you to request a quota increase. For example, ROSA associates one security group with network interfaces in worker machine pools by default. Because the default quota for security groups per

network interface is 5, if you want to add 5 custom security groups you must request a quota increase to bring the total number of security groups on worker network interfaces to 6.

To increase quotas, open the Service Quotas console and request a quota increase. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

> **Note**
>
> You can check your quotas using the AWS SDKs, but the SDK calculation doesn't include existing ROSA resources. The quota check in the SDK may pass, and ROSA cluster creation may fail. To fix this issue, open the Service Quotas console and request a quota increase.

**ROSA-required quotas**

| Name | Service code | Default | Minimum required | Adjustable | Description |
|---|---|---|---|---|---|
| Running On-Demand Standard (A, C, D, H, I, M, R, T, Z) instances | ec2 | 5 | 100 | Yes | Maximum number of vCPUs assigned to the Running On-Demand Standard (A, C, D, H, I, M, R, T, Z) instances.<br><br>The default value of 5 vCPUs is not sufficient to create ROSA clusters. ROSA requires 100 vCPUs |

| Name | Service code | Default | Minimum required | Adjustable | Description |
|---|---|---|---|---|---|
| | | | | | for cluster creation. |
| Storage for General Purpose SSD (gp3) volumes, in TiB | ebs | 50 | 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across General Purpose SSD (gp3) volumes in this Region. |
| Storage for General Purpose SSD (gp2) volumes, in TiB | ebs | 50 | 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across General Purpose SSD (gp2) volumes in this Region. |

| Name | Service code | Default | Minimum required | Adjustable | Description |
|---|---|---|---|---|---|
| Storage for Provisioned IOPS SSD (io1) volumes, in TiB | ebs | 50 | 300 | Yes | The maximum aggregated amount of storage, in TiB, that can be provisioned across Provisioned IOPS SSD (io1) volumes in this Region. |

**Note**

The default values are the initial quotas set by AWS. These default values are separate from the actual applied quota values and maximum possible service quotas. For more information, see Terminology in Service Quotas in the *Service Quotas User Guide*.

## General AWS service quotas

ROSA uses the following default quotas for Amazon EC2, Amazon VPC, Amazon EBS, and Elastic Load Balancing.

### Amazon EC2

- EC2-VPC Elastic IPs

### Amazon VPC

- VPCs per Region
- Network interfaces per Region

- Internet gateways per Region

**Amazon EBS**

- Snapshots per Region
- IOPS for Provisioned IOPS SSD (io1) volumes

**Elastic Load Balancing**

- Application Load Balancers per Region
- Classic Load Balancers per Region

# Amazon Route 53 endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Hosted zones, records, health checks, DNS query logs, reusable delegation sets, traffic policies, and cost allocation tags

When you use the **AWS CLI or SDKs** to submit requests, you can either leave the Region and endpoint unspecified, or specify the applicable Region:

- Route 53 in AWS Regions other than the Beijing and Ningxia Regions: specify us-east-1 as the Region.
- Route 53 in the Beijing and Ningxia Regions: specify cn-northwest-1.

When you use the **Route 53 API** to submit requests, use the same Regions as above to sign requests. The IPv4 service endpoint is route53.amazonaws.com and the dual-stack service endpoint is route53.global.api.aws.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | route53.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | route53.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | route53.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | route53.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | route53.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | route53.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | route53.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | route53.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | route53.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | route53.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | route53.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | route53.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | route53.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | route53.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Paris) | eu-west-3 | route53.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | route53.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | route53.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | route53.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | route53.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | route53.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | route53.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | route53.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | route53.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | route53.us-gov.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | route53.us-gov.amazonaws.com | HTTPS | |

## Requests for domain registration

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | route53domains.us-east-1.amazonaws.com | HTTPS | |

## Requests for Route 53 Resolver

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | route53resolver.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | route53resolver.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | route53resolver.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | route53resolver.us-west-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | route53resolver.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | route53resolver.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | route53resolver.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | route53resolver.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | route53resolver.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | route53resolver.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | route53resolver.ap-south-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | route53resolver.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | route53resolver.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | route53resolver.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | route53resolver.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | route53resolver.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | route53resolver.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | route53resolver.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | route53resolver.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | route53resolver.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | route53resolver.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | route53resolver.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | route53resolver.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | route53resolver.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | route53resolver.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | route53resolver.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | route53resolver.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | route53resolver.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | route53resolver.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | route53resolver.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | route53resolver.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | route53resolver.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | route53resolver.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | route53resolver.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | route53resolver.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | route53resolver.us-gov-west-1.amazonaws.com | HTTPS | |

## Requests for Route 53 Global Resolver

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | route53globalresolver.us-east-2.api.aws<br>route53globalresolver-fips.us-east-2.api.aws | HTTPS<br>HTTPS |

## Requests for Route 53 Profiles

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | route53profiles.us-east-2.amazonaws.com | HTTPS |
| | | route53profiles-fips.us-east-2.api.aws | HTTPS |
| | | route53profiles.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | route53profiles.us-east-1.amazonaws.com | HTTPS |
| | | route53profiles-fips.us-east-1.api.aws | HTTPS |
| | | route53profiles.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | route53profiles.us-west-1.amazonaws.com | HTTPS |
| | | route53profiles-fips.us-west-1.api.aws | HTTPS |
| | | route53profiles.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | route53profiles.us-west-2.amazonaws.com | HTTPS |
| | | route53profiles-fips.us-west-2.api.aws | HTTPS |
| | | route53profiles.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | route53profiles.af-south-1.amazonaws.com | HTTPS |
| | | route53profiles.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | route53profiles.ap-east-1.amazonaws.com | HTTPS |
| | | route53profiles.ap-east-1.api.aws | HTTPS |
| Asia Pacific | ap-south-2 | route53profiles.ap-south-2.amazonaws.com | HTTPS |
| | | route53profiles.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | route53profiles.ap-southeast-3.amazonaws.com<br><br>route53profiles.ap-southeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | route53profiles.ap-southeast-5.amazonaws.com<br><br>route53profiles.ap-southeast-5.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | route53profiles.ap-southeast-4.amazonaws.com<br><br>route53profiles.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | route53profiles.ap-south-1.amazonaws.com<br><br>route53profiles.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | route53profiles.ap-southeast-6.amazonaws.com<br><br>route53profiles.ap-southeast-6.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | route53profiles.ap-northeast-3.amazonaws.com<br><br>route53profiles.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | route53profiles.ap-northeast-2.amazonaws.com<br><br>route53profiles.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | route53profiles.ap-southeast-1.amazonaws.com<br>route53profiles.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | route53profiles.ap-southeast-2.amazonaws.com<br>route53profiles.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | route53profiles.ap-east-2.amazonaws.com<br>route53profiles.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | route53profiles.ap-southeast-7.amazonaws.com<br>route53profiles.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | route53profiles.ap-northeast-1.amazonaws.com<br>route53profiles.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | route53profiles.ca-central-1.amazonaws.com<br>route53profiles-fips.ca-central-1.api.aws<br>route53profiles.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | route53profiles.ca-west-1.amazonaws.com<br>route53profiles.ca-west-1.api.aws<br>route53profiles-fips.ca-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | route53profiles.eu-central-1.amazonaws.com<br>route53profiles.eu-central-1.api.aws | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | route53profiles.eu-west-1.amazonaws.com<br>route53profiles.eu-west-1.api.aws | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | route53profiles.eu-west-2.amazonaws.com<br>route53profiles.eu-west-2.api.aws | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | route53profiles.eu-south-1.amazonaws.com<br>route53profiles.eu-south-1.api.aws | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | route53profiles.eu-west-3.amazonaws.com<br>route53profiles.eu-west-3.api.aws | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | route53profiles.eu-south-2.amazonaws.com<br>route53profiles.eu-south-2.api.aws | HTTPS<br>HTTPS |
| Europe (Stockholm) | eu-north-1 | route53profiles.eu-north-1.amazonaws.com<br>route53profiles.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | route53profiles.eu-central-2.amazonaws.com<br>route53profiles.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | route53profiles.il-central-1.amazonaws.com<br>route53profiles.il-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | route53profiles.mx-central-1.amazonaws.com<br>route53profiles.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | route53profiles.me-south-1.amazonaws.com<br>route53profiles.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | route53profiles.me-central-1.amazonaws.com<br>route53profiles.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | route53profiles.sa-east-1.amazonaws.com<br>route53profiles.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | route53profiles.us-gov-east-1.amazonaws.com<br>route53profiles.us-gov-east-1.api.aws<br>route53profiles-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | route53profiles.us-gov-west-1.amazonaws.com<br>route53profiles-fips.us-gov-west-1.api.aws<br>route53profiles.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

## Requests for Route 53 auto naming

Amazon Route 53 auto naming has been released as a separate service, AWS Cloud Map. For a list of service endpoints, see Service endpoints. For AWS Cloud Map documentation, see AWS Cloud Map Documentation.

# Service quotas

The following are the service quotas for Route 53

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Amazon VPCs that you can associate with a private hosted zone | Each supported Region: 300 | Yes | The maximum number of Amazon VPCs that you can associate with a private hosted zone |
| CIDR blocks per collection | Each supported Region: 1,000 | Yes | The maximum number of CIDR blocks that you can create per CIDR collection |
| CIDR collections | Each supported Region: 5 | Yes | The maximum number of CIDR collections that you can create per account |
| Child health checks that a calculated health check can monitor | Each supported Region: 255 | No | The maximum number of child health checks that a calculated health check can monitor |
| Domain count limit | Each supported Region: 20 | Yes | The maximum number of domains you can register using this account |
| Health checks | Each supported Region: 200 | Yes | The maximum number of health checks that you can create using this account |
| Hosted zones | Each supported Region: 500 | Yes | The maximum number of hosted zones that |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | you can create using this account |
| Hosted zones that can use the same reusable delegation set | Each supported Region: 100 | Yes | The maximum number of hosted zones that can use the same reusable delegation set |
| Records per hosted zone | Each supported Region: 10,000 | Yes | The maximum number of records that you can create in a hosted zone |
| Reusable delegation sets | Each supported Region: 100 | Yes | The maximum number of reusable delegation sets that you can create using this account |
| Traffic flow policies | Each supported Region: 50 | Yes | The maximum number of traffic flow policies that you can create using this account |
| Traffic flow policy records | Each supported Region: 5 | Yes | The maximum number of traffic flow policy records that you can create using this account |

The following quotas are for Route 53 Resolver.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Associations between resolver rules and VPCs per AWS Region | Each supported Region: 2,000 | Yes | Maximum number of associations between |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | resolver rules and VPCs per AWS Region |
| DNS Firewall rule group associations per VPC | Each supported Region: 5 | No | The maximum number of DNS Firewall rule groups that you can associate to a VPC. |
| DNS Firewall rule groups associations to a Route 53 Profile. | Each supported Region: 5 | No | The maximum number of DNS Firewall rule groups that can be associated to a Route 53 Profile. |
| DNS Firewall rules groups per Region | Each supported Region: 1,000 | Yes | The maximum number of DNS Firewall rules groups per Region. |
| Domain lists per account | Each supported Region: 1,000 | Yes | The maximum number of domain lists for an account. |
| Domains in a file imported from S3 | Each supported Region: 250,000 | Yes | The maximum number of domains that you can import from a single file thats stored in an Amazon S3 bucket. |
| Domains per account | Each supported Region: 100,000 | Yes | The maximum number of domains that you can specify across all of the domain lists for an account. |
| IP addresses per resolver endpoint | Each supported Region: 6 | Yes | Maximum number of IP addresses per resolver endpoint |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of resolver endpoints per AWS Region | Each supported Region: 4 | Yes | Resolver endpoints per AWS Region |
| Resolver rule associations to a Route 53 Profile. | Each supported Region: 1,000 | Yes | The maximum number of Resolver rule associations per Route 53 Profile. |
| Resolver rules per AWS Region | Each supported Region: 1,000 | Yes | Maximum number of resolver rules per AWS Region |
| Rules in a DNS Firewall rule group | Each supported Region: 100 | Yes | The maximum number of rules in a DNS Firewall rule group. |
| Target IP addresses per resolver rule | Each supported Region: 6 | No | Maximum number of target IP addresses per resolver rule |

The following quotas are for Route 53 Profiles

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Private hosted zone associations to a Route 53 Profile. | Each supported Region: 5,000 | Yes | The maximum number of private hosted zones that can be associated to a Route 53 Profile. |
| Route 53 Profiles per account per Region. | Each supported Region: 5 | Yes | The maximum number of Route 53 Profiles per AWS account per Region. |
| VPC Endpoint associations to a Route 53 Profile. | Each supported Region: 100 | Yes | The maximum number of VPC Endpoints that can |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| | | | be associated to a Route 53 Profile. |
| VPC to Route 53 Profile associations per Region | Each supported Region: 1,000 | Yes | The maximum number of VPCs that can be associated to a Route 53 Profile. |

For more information, see Route 53 quotas in the *Amazon Route 53 Developer Guide*.

# Amazon SageMaker AI endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service Endpoints

The following table provides a list of Region-specific endpoints that Amazon SageMaker AI supports for training and deploying models. This include creating and managing notebook instances, training jobs, model, endpoint configurations, and endpoints.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | api.sagemaker.us-east-2.amazonaws.com | HTTPS | |
| | | api-fips.sagemaker.us-east-2.api.aws | HTTPS | |
| | | api-fips.sagemaker.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | api.sagemaker.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | api.sagemaker.us-east-1.amazonaws.com | HTTPS |
| | | api-fips.sagemaker.us-east-1.api.aws | HTTPS |
| | | api-fips.sagemaker.us-east-1.amazonaws.com | HTTPS |
| | | api.sagemaker.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | api.sagemaker.us-west-1.amazonaws.com | HTTPS |
| | | api-fips.sagemaker.us-west-1.api.aws | HTTPS |
| | | api-fips.sagemaker.us-west-1.amazonaws.com | HTTPS |
| | | api.sagemaker.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | api.sagemaker.us-west-2.amazonaws.com | HTTPS |
| | | api-fips.sagemaker.us-west-2.api.aws | HTTPS |
| | | api-fips.sagemaker.us-west-2.amazonaws.com | HTTPS |
| | | api.sagemaker.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | api.sagemaker.af-south-1.amazonaws.com | HTTPS |
| | | api.sagemaker.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | api.sagemaker.ap-east-1.amazonaws.com | HTTPS |
| | | api.sagemaker.ap-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | api.sagemaker.ap-south-2.amazonaws.com<br>api.sagemaker.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | api.sagemaker.ap-southeast-3.amazonaws.com<br>api.sagemaker.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | api.sagemaker.ap-southeast-5.amazonaws.com<br>api.sagemaker.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | api.sagemaker.ap-southeast-4.amazonaws.com<br>api.sagemaker.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | api.sagemaker.ap-south-1.amazonaws.com<br>api.sagemaker.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | api.sagemaker.ap-southeast-6.amazonaws.com<br>api.sagemaker.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | api.sagemaker.ap-northeast-3.amazonaws.com<br>api.sagemaker.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | api.sagemaker.ap-northeast-2.amazonaws.com<br>api.sagemaker.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | api.sagemaker.ap-southeast-1.amazonaws.com<br>api.sagemaker.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | api.sagemaker.ap-southeast-2.amazonaws.com<br>api.sagemaker.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | api.sagemaker.ap-east-2.amazonaws.com<br>api.sagemaker.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | api.sagemaker.ap-southeast-7.amazonaws.com<br>api.sagemaker.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | api.sagemaker.ap-northeast-1.amazonaws.com<br>api.sagemaker.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | api.sagemaker.ca-central-1.amazonaws.com<br>api-fips.sagemaker.ca-central-1.api.aws<br>api-fips.sagemaker.ca-central-1.amazonaws.com<br>api.sagemaker.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | api.sagemaker.ca-west-1.amazonaws.com | HTTPS | |
| | | api-fips.sagemaker.ca-west-1.api.aws | HTTPS | |
| | | api-fips.sagemaker.ca-west-1.amazonaws.com | HTTPS | |
| | | api.sagemaker.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | api.sagemaker.eu-central-1.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | api.sagemaker.eu-west-1.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | api.sagemaker.eu-west-2.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | api.sagemaker.eu-south-1.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | api.sagemaker.eu-west-3.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | api.sagemaker.eu-south-2.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | api.sagemaker.eu-north-1.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | api.sagemaker.eu-central-2.amazonaws.com | HTTPS | |
| | | api.sagemaker.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | api.sagemaker.il-central-1.amazonaws.com<br>api.sagemaker.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | api.sagemaker.mx-central-1.amazonaws.com<br>api.sagemaker.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | api.sagemaker.me-south-1.amazonaws.com<br>api.sagemaker.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | api.sagemaker.me-central-1.amazonaws.com<br>api.sagemaker.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | api.sagemaker.sa-east-1.amazonaws.com<br>api.sagemaker.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | api.sagemaker.us-gov-east-1.amazonaws.com<br>api-fips.sagemaker.us-gov-east-1.api.aws<br>api-fips.sagemaker.us-gov-east-1.amazonaws.com<br>api.sagemaker.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | api.sagemaker.us-gov-west-1.amazonaws.com<br>api-fips.sagemaker.us-gov-west-1.api.aws<br>api-fips.sagemaker.us-gov-west-1.amazonaws.com<br>api.sagemaker.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

The following table provides a list of Region-specific endpoints that Amazon SageMaker AI supports for making inference requests against models hosted in SageMaker AI.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | runtime.sagemaker.us-east-2.amazonaws.com | HTTPS | |
| | | runtime-fips.sagemaker.us-east-2.api.aws | HTTPS | |
| | | runtime-fips.sagemaker.us-east-2.amazonaws.com | HTTPS | |
| | | runtime.sagemaker.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | runtime.sagemaker.us-east-1.amazonaws.com | HTTPS | |
| | | runtime-fips.sagemaker.us-east-1.api.aws | HTTPS | |
| | | runtime-fips.sagemaker.us-east-1.amazonaws.com | HTTPS | |
| | | runtime.sagemaker.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | runtime.sagemaker.us-west-1.amazonaws.com | HTTPS | |
| | | runtime-fips.sagemaker.us-west-1.api.aws | HTTPS | |
| | | runtime-fips.sagemaker.us-west-1.amazonaws.com | HTTPS | |
| | | runtime.sagemaker.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | runtime.sagemaker.us-west-2.amazonaws.com | HTTPS | |
| | | runtime-fips.sagemaker.us-west-2.api.aws | HTTPS | |
| | | runtime-fips.sagemaker.us-west-2.amazonaws.com | HTTPS | |
| | | runtime.sagemaker.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | runtime.sagemaker.af-south-1.amazonaws.com<br>runtime.sagemaker.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | runtime.sagemaker.ap-east-1.amazonaws.com<br>runtime.sagemaker.ap-east-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | runtime.sagemaker.ap-south-2.amazonaws.com<br>runtime.sagemaker.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | runtime.sagemaker.ap-southeast-3.amazonaws.com<br>runtime.sagemaker.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | runtime.sagemaker.ap-southeast-5.amazonaws.com<br>runtime.sagemaker.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | runtime.sagemaker.ap-southeast-4.amazonaws.com<br>runtime.sagemaker.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | runtime.sagemaker.ap-south-1.amazonaws.com<br>runtime.sagemaker.ap-south-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (New Zealand) | ap-southeast-6 | runtime.sagemaker.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | runtime.sagemaker.ap-northeast-3.amazonaws.com<br>runtime.sagemaker.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | runtime.sagemaker.ap-northeast-2.amazonaws.com<br>runtime.sagemaker.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | runtime.sagemaker.ap-southeast-1.amazonaws.com<br>runtime.sagemaker.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | runtime.sagemaker.ap-southeast-2.amazonaws.com<br>runtime.sagemaker.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | runtime.sagemaker.ap-east-2.amazonaws.com<br>runtime.sagemaker.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | runtime.sagemaker.ap-southeast-7.amazonaws.com<br>runtime.sagemaker.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | runtime.sagemaker.ap-northeast-1.amazonaws.com<br>runtime.sagemaker.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | runtime.sagemaker.ca-central-1.amazonaws.com<br>runtime.sagemaker.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | runtime.sagemaker.ca-west-1.amazonaws.com<br>runtime.sagemaker.ca-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | runtime.sagemaker.eu-central-1.amazonaws.com<br>runtime.sagemaker.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | runtime.sagemaker.eu-west-1.amazonaws.com<br>runtime.sagemaker.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | runtime.sagemaker.eu-west-2.amazonaws.com<br>runtime.sagemaker.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | runtime.sagemaker.eu-south-1.amazonaws.com<br>runtime.sagemaker.eu-south-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | runtime.sagemaker.eu-west-3.amazonaws.com<br>runtime.sagemaker.eu-west-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | runtime.sagemaker.eu-south-2.amazonaws.com<br><br>runtime.sagemaker.eu-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | runtime.sagemaker.eu-north-1.amazonaws.com<br><br>runtime.sagemaker.eu-north-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Zurich) | eu-central-2 | runtime.sagemaker.eu-central-2.amazonaws.com<br><br>runtime.sagemaker.eu-central-2.api.aws | HTTPS<br><br>HTTPS | |
| Israel (Tel Aviv) | il-central-1 | runtime.sagemaker.il-central-1.amazonaws.com<br><br>runtime.sagemaker.il-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Mexico (Central) | mx-central-1 | runtime.sagemaker.mx-central-1.amazonaws.com<br><br>runtime.sagemaker.mx-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | runtime.sagemaker.me-south-1.amazonaws.com<br><br>runtime.sagemaker.me-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (UAE) | me-central-1 | runtime.sagemaker.me-central-1.amazonaws.com<br><br>runtime.sagemaker.me-central-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | runtime.sagemaker.sa-east-1.amazonaws.com<br>runtime.sagemaker.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | runtime.sagemaker.us-gov-east-1.amazonaws.com<br>runtime.sagemaker.us-gov-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | runtime.sagemaker.us-gov-west-1.amazonaws.com<br>runtime.sagemaker.us-gov-west-1.api.aws | HTTPS<br>HTTPS | |

The following table provides a list of Region-specific endpoints that Amazon SageMaker AI supports for SageMaker AI Feature Store.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | featurestore-runtime.sagemaker.us-east-2.amazonaws.com<br>featurestore-runtime-fips.sagemaker.us-east-2.api.aws<br>featurestore-runtime-fips.sagemaker.us-east-2.amazonaws.com<br>featurestore-runtime.sagemaker.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | featurestore-runtime.sagemaker.us-east-1.amazonaws.com<br>featurestore-runtime-fips.sagemaker.us-east-1.api.aws<br>featurestore-runtime-fips.sagemaker.us-east-1.amazonaws.com<br>featurestore-runtime.sagemaker.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | featurestore-runtime.sagemaker.us-west-1.amazonaws.com<br>featurestore-runtime-fips.sagemaker.us-west-1.api.aws<br>featurestore-runtime-fips.sagemaker.us-west-1.amazonaws.com<br>featurestore-runtime.sagemaker.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | featurestore-runtime.sagemaker.us-west-2.amazonaws.com<br>featurestore-runtime-fips.sagemaker.us-west-2.api.aws<br>featurestore-runtime-fips.sagemaker.us-west-2.amazonaws.com<br>featurestore-runtime.sagemaker.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | featurestore-runtime.sagemaker.af-south-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.af-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | featurestore-runtime.sagemaker.ap-east-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-east-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | featurestore-runtime.sagemaker.ap-southeast-3.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-southeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | featurestore-runtime.sagemaker.ap-south-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | featurestore-runtime.sagemaker.ap-northeast-3.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | featurestore-runtime.sagemaker.ap-northeast-2.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | featurestore-runtime.sagemaker.ap-southeast-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | featurestore-runtime.sagemaker.ap-southeast-2.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | featurestore-runtime.sagemaker.ap-northeast-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Canada (Central) | ca-central-1 | featurestore-runtime.sagemaker.ca-central-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.ca-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | featurestore-runtime.sagemaker.eu-central-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | featurestore-runtime.sagemaker.eu-west-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | featurestore-runtime.sagemaker.eu-west-2.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-west-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Milan) | eu-south-1 | featurestore-runtime.sagemaker.eu-south-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Paris) | eu-west-3 | featurestore-runtime.sagemaker.eu-west-3.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-west-3.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Spain) | eu-south-2 | featurestore-runtime.sagemaker.eu-south-2.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-south-2.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | featurestore-runtime.sagemaker.eu-north-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-north-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Zurich) | eu-central-2 | featurestore-runtime.sagemaker.eu-central-2.amazonaws.com<br><br>featurestore-runtime.sagemaker.eu-central-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | featurestore-runtime.sagemaker.il-central-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.il-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Middle East (Bahrain) | me-south-1 | featurestore-runtime.sagemaker.me-south-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.me-south-1.api.aws | HTTPS<br><br>HTTPS | |
| South America (São Paulo) | sa-east-1 | featurestore-runtime.sagemaker.sa-east-1.amazonaws.com<br><br>featurestore-runtime.sagemaker.sa-east-1.api.aws | HTTPS<br><br>HTTPS | |

## Service Quotas

Depending on your activities and resource usage over time, your Amazon SageMaker AI quotas might be different from the default SageMaker AI quotas listed in the following table. The default quotas in this page are based on new accounts. If you encounter an error message that you've exceeded your quota, use the Service Quotas console to request an increase. For instructions on how to request a quota increase, see Requesting a quota increase.

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Canvas Apps running on ml.m5.4xlarge instances | Each supported Region: 10 | Yes | Canvas Apps running on ml.m5.4xlarge instances |
| Canvas Apps running on system instances | Each supported Region: 10 | Yes | Canvas Apps running on system instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Large-sized MLflow Tracking Server usage | Each supported Region: 1 | Yes | Large-sized MLflow Tracking Server usage |
| Longest run time for a processing job | Each supported Region: 432,000 Seconds | Yes | Longest run time for a processing job |
| Longest run time for a training job | Each supported Region: 432,000 | Yes | Longest run time for a training job |
| Longest run time for an AutoML job, from creation to termination | Each supported Region: 2,592,000 Seconds | No | Longest run time for an AutoML job, from creation to termination |
| Maximum dataset size AutoML job can be run on | Each supported Region: 100 | No | Maximum dataset size AutoML job can be run on |
| Maximum number Spot instances allowed per SageMaker HyperPod cluster | Each supported Region: 20 | Yes | Maximum number Spot instances allowed per SageMaker HyperPod cluster |
| Maximum number instances allowed per SageMaker HyperPod cluster | Each supported Region: 20 | Yes | Maximum number instances allowed per SageMaker HyperPod cluster |
| Maximum number of A2I flow definitions | Each supported Region: 100 | Yes | Maximum number of A2I flow definitions |
| Maximum number of A2I human task UIs | Each supported Region: 100 | Yes | Maximum number of A2I human task UIs |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of Ground Truth Streaming labeling jobs | Each supported Region: 0 | Yes | Maximum number of Ground Truth Streaming labeling jobs |
| Maximum number of Ground Truth labeling jobs | Each supported Region: 1 | Yes | Maximum number of Ground Truth labeling jobs |
| Maximum number of MLflow Tracking Servers | Each supported Region: 10 | Yes | Maximum number of MLflow Tracking Servers |
| Maximum number of SageMaker Model Package Groups allowed per account | Each supported Region: 250 | Yes | Maximum number of SageMaker Model Package Groups allowed per account |
| Maximum number of SageMaker Model Package allowed per account | Each supported Region: 10,000 | Yes | Maximum number of SageMaker Model Package allowed per account |
| Maximum number of SageMaker Projects allowed per account | Each supported Region: 500 | Yes | Maximum number of SageMaker Projects allowed per account |
| Maximum number of SageMakerImage images allowed per account | Each supported Region: 250 | Yes | Maximum number of SageMakerImage images allowed per account |
| Maximum number of Studio spaces allowed per account | Each supported Region: 6 | Yes | Maximum number of Studio spaces allowed per account |
| Maximum number of Studio user profiles allowed per account | Each supported Region: 2 | Yes | Maximum number of Studio user profiles allowed per account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of cluster scheduler config versions allowed per account in each region. | Each supported Region: 100 | Yes | The maximum number of cluster scheduler config versions that you can create in this account in the current Region. Cluster scheduler config policy could be used for tasks governance in SageMaker. |
| Maximum number of compute quota versions allowed per account in each region. | Each supported Region: 100 | Yes | The maximum number of compute quota versions that you can create in this account in the current Region. Compute quota policy could be used for tasks governance in SageMaker. |
| Maximum number of concurrent AutoML Jobs | Each supported Region: 1 | Yes | Maximum number of concurrent AutoML Jobs |
| Maximum number of concurrent model customization serverless jobs per Region | Each supported Region: 0 | Yes | Maximum number of concurrent model customization serverless jobs per Region |
| Maximum number of concurrent model evaluation serverless jobs per Region | Each supported Region: 0 | Yes | Maximum number of concurrent model evaluation serverless jobs per Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of concurrent multi-turn reinforcement evaluation jobs | Each supported Region: 0 | Yes | Maximum number of concurrent multi-turn reinforcement evaluation jobs |
| Maximum number of concurrent multi-turn reinforcement fine tuning jobs | Each supported Region: 0 | Yes | Maximum number of concurrent multi-turn reinforcement fine tuning jobs |
| Maximum number of concurrent pipeline executions allowed per account | Each supported Region: 20 | Yes | Maximum number of concurrent pipeline executions allowed per account |
| Maximum number of concurrently running model card export jobs allowed per account. | Each supported Region: 20 | Yes | Maximum number of concurrently running model card export jobs allowed per account. |
| Maximum number of dataset objects per labeling job | Each supported Region: 10,000 | No | Maximum number of dataset objects per labeling job |
| Maximum number of hub contents per private hub | Each supported Region: 10,000 | Yes | Maximum number of hub contents per private hub |
| Maximum number of hyper parameter tuning jobs that can run at once in parallel | Each supported Region: 100 | No | Maximum number of hyper parameter tuning jobs that can run at once in parallel |
| Maximum number of instances per endpoint | Each supported Region: 4 | Yes | Maximum number of instances per endpoint |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of instances per processing job | Each supported Region: 20 | Yes | Maximum number of instances per processing job |
| Maximum number of instances per spot training job | Each supported Region: 20 | Yes | Maximum number of instances per spot training job |
| Maximum number of instances per training job | Each supported Region: 20 | Yes | Maximum number of instances per training job |
| Maximum number of model card versions allowed per account. | Each supported Region: 100 | Yes | Maximum number of model card versions allowed per account. |
| Maximum number of parallel AI benchmark jobs | Each supported Region: 2 | Yes | Maximum number of parallel AI benchmark jobs |
| Maximum number of parallel AI recommendation jobs | Each supported Region: 2 | Yes | Maximum number of parallel AI recommendation jobs |
| Maximum number of parallel compilation jobs | Each supported Region: 2 | Yes | Maximum number of parallel compilation jobs |
| Maximum number of parameters allowed per pipeline | Each supported Region: 200 | No | Maximum number of parameters allowed per pipeline |
| Maximum number of pipelines allowed per account | Each supported Region: 500 | Yes | Maximum number of pipelines allowed per account |
| Maximum number of private hubs per account | Each supported Region: 100 | Yes | Maximum number of private hubs per account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of running Studio apps allowed per account | Each supported Region: 40 | Yes | Maximum number of running Studio apps allowed per account |
| Maximum number of serverless endpoints | Each supported Region: 5 | Yes | Maximum number of serverless endpoints |
| Maximum number of steps allowed per pipeline | Each supported Region: 50 | No | Maximum number of steps allowed per pipeline |
| Maximum number of training jobs each hyper parameter tuning job can run in parallel at once | Each supported Region: 10 | No | Maximum number of training jobs each hyper parameter tuning job can run in parallel at once |
| Maximum number of training jobs that each hyperparameter tuning job can create | Each supported Region: 750 | No | Maximum number of training jobs that each hyperparameter tuning job can create |
| Maximum number of training jobs that each hyperparameter tuning job with Random search strategy can create | Each supported Region: 500 | Yes | Maximum number of training jobs that each hyperparameter tuning job with Random search strategy can create |
| Maximum size of EBS volume in GB for a SageMaker HyperPod cluster instance | Each supported Region: 0 | Yes | Maximum size of EBS volume in GB for a SageMaker HyperPod cluster instance |
| Maximum subsampled dataset size AutoML job can be run on | Each supported Region: 5 | No | Maximum subsampled dataset size AutoML job can be run on |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum total FSx Lustre capacity in TiB across SageMaker HyperPod restricted instance groups | Each supported Region: 0 | Yes | Maximum total FSx Lustre capacity in TiB across SageMaker HyperPod restricted instance groups |
| Maximum total concurrency that can be allocated across all serverless endpoints | Each supported Region: 10 | Yes | Maximum total concurrency that can be allocated across all serverless endpoints |
| Medium-sized MLflow Tracking Server usage | Each supported Region: 5 | Yes | Medium-sized MLflow Tracking Server usage |
| Number of elastic inference accelerators across active endpoints | Each supported Region: 0 | Yes | Number of elastic inference accelerators across active endpoints |
| Number of instances across active endpoints | Each supported Region: 4 | Yes | Number of instances across active endpoints |
| Number of instances across all processing jobs | Each supported Region: 4 | Yes | Number of instances across all processing jobs |
| Number of instances across all spot training jobs | Each supported Region: 4 | Yes | Number of instances across all spot training jobs |
| Number of instances across all training jobs | Each supported Region: 4 | Yes | Number of instances across all training jobs |
| Number of instances across all transform jobs | Each supported Region: 4 | Yes | Number of instances across all transform jobs |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ml.p4d.24xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p4d.24xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |
| Number of ml.p4de.24xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p4de.24xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ml.p5.48xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p5.48xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |
| Number of ml.p5.4xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p5.4xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ml.p5e.48xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p5e.48xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |
| Number of ml.p5en.48xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p5en.48xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ml.p6-b200.48xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p6-b200.48xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |
| Number of ml.p6-b300.48xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p6-b300.48xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ml.p6e-gb200.36xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.p6e-gb200.36xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |
| Number of ml.trn1.32xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.trn1.32xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of ml.trn2.48xlarge instances in reserved capacity across training plans per Region | Each supported Region: 0 | Yes | The maximum number of ml.trn2.48xlarge instances available in SageMaker reserved capacities across all training plans in this account and Region at any given time. These reservations secure compute resources across multiple Availability Zones, ensuring availability for large-scale AI model training workloads. |
| Number of training plans per Region | Each supported Region: 25 | Yes | The maximum number of training plans that you can create in this account in the current Region. A training plan allows you to reserve and manage GPU capacity for large-scale AI model training within specified timelines and budgets, offering predictable access to resources for SageMaker training jobs or Hyperpod clusters. |
| Number of workteams | Each supported Region: 25 | Yes | Number of workteams |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RSessionGateway Apps running on ml.c5.12xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.c5.12x large instance |
| RSessionGateway Apps running on ml.c5.18xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.c5.18x large instance |
| RSessionGateway Apps running on ml.c5.24xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.c5.24x large instance |
| RSessionGateway Apps running on ml.c5.2xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.c5.2xlarge instance |
| RSessionGateway Apps running on ml.c5.4xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.c5.4xlarge instance |
| RSessionGateway Apps running on ml.c5.9xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.c5.9xlarge instance |
| RSessionGateway Apps running on ml.c5.large instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.c5.large instance |
| RSessionGateway Apps running on ml.c5.xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.c5.xlarge instance |
| RSessionGateway Apps running on ml.g4dn.12xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.g4dn.12xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RSessionGateway Apps running on ml.g4dn.16xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.g4dn.16xlarge instance |
| RSessionGateway Apps running on ml.g4dn.2xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.g4dn.2xlarge instance |
| RSessionGateway Apps running on ml.g4dn.4xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.g4dn.4xlarge instance |
| RSessionGateway Apps running on ml.g4dn.8xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.g4dn.8xlarge instance |
| RSessionGateway Apps running on ml.g4dn.xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.g4dn.xlarge instance |
| RSessionGateway Apps running on ml.m5.12xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5.12xlarge instance |
| RSessionGateway Apps running on ml.m5.16xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5.16xlarge instance |
| RSessionGateway Apps running on ml.m5.24xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5.24xlarge instance |
| RSessionGateway Apps running on ml.m5.2xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5.2xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RSessionGateway Apps running on ml.m5.4xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5.4xlarge instance |
| RSessionGateway Apps running on ml.m5.8xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5.8xlarge instance |
| RSessionGateway Apps running on ml.m5.large instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5.large instance |
| RSessionGateway Apps running on ml.m5.xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5.xlarge instance |
| RSessionGateway Apps running on ml.m5d.12xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5d.12xlarge instance |
| RSessionGateway Apps running on ml.m5d.16xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5d.16xlarge instance |
| RSessionGateway Apps running on ml.m5d.24xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5d.24xlarge instance |
| RSessionGateway Apps running on ml.m5d.2xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5d.2xlarge instance |
| RSessionGateway Apps running on ml.m5d.4xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5d.4xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RSessionGateway Apps running on ml.m5d.8xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.m5d.8xlarge instance |
| RSessionGateway Apps running on ml.m5d.large instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5d.large instance |
| RSessionGateway Apps running on ml.m5d.xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.m5d.xlarge instance |
| RSessionGateway Apps running on ml.p3.16xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.p3.16xlarge instance |
| RSessionGateway Apps running on ml.p3.2xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.p3.2xlarge instance |
| RSessionGateway Apps running on ml.p3.8xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.p3.8xlarge instance |
| RSessionGateway Apps running on ml.p3dn.24xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.p3dn.24xlarge instance |
| RSessionGateway Apps running on ml.p4d.24xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.p4d.24xlarge instance |
| RSessionGateway Apps running on ml.r5.12xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.r5.12xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RSessionGateway Apps running on ml.r5.16xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.r5.16xlarge instance |
| RSessionGateway Apps running on ml.r5.24xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.r5.24xlarge instance |
| RSessionGateway Apps running on ml.r5.2xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.r5.2xlarge instance |
| RSessionGateway Apps running on ml.r5.4xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.r5.4xlarge instance |
| RSessionGateway Apps running on ml.r5.8xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.r5.8xlarge instance |
| RSessionGateway Apps running on ml.r5.large instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.r5.large instance |
| RSessionGateway Apps running on ml.r5.xlarge instance | Each supported Region: 1 | Yes | RSessionGateway Apps running on ml.r5.xlarge instance |
| RSessionGateway Apps running on ml.t3.2xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.t3.2xlarge instance |
| RSessionGateway Apps running on ml.t3.large instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.t3.large instance |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| RSessionGateway Apps running on ml.t3.medium instance | Each supported Region: 2 | Yes | RSessionGateway Apps running on ml.t3.medium instance |
| RSessionGateway Apps running on ml.t3.xlarge instance | Each supported Region: 0 | Yes | RSessionGateway Apps running on ml.t3.xlarge instance |
| RStudioServerPro Apps running on ml.c5.4xlarge instances | Each supported Region: 0 | Yes | RStudioServerPro Apps running on ml.c5.4xlarge instances |
| RStudioServerPro Apps running on ml.c5.9xlarge instances | Each supported Region: 0 | Yes | RStudioServerPro Apps running on ml.c5.9xlarge instances |
| RStudioServerPro Apps running on system instances | Each supported Region: 1 | Yes | RStudioServerPro Apps running on system instances |
| Rate of CreateEndpoint requests | Each supported Region: 1 | No | The maximum number of CreateEndpoint requests that you can make, per second, in this account in the current region. |
| Rate of CreateEndpointConfig requests | Each supported Region: 1 | No | The maximum number of CreateEndpointConfig requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateFlowDefinition requests | Each supported Region: 1 | No | The maximum number of CreateFlowDefinition requests that you can make, per second, in this account in the current region. |
| Rate of CreateHumanTaskUi requests | Each supported Region: 1 | No | The maximum number of CreateHumanTaskUi requests that you can make, per second, in this account in the current region. |
| Rate of CreateLabelingJob requests | Each supported Region: 1 | No | The maximum number of CreateLabelingJob requests that you can make, per second, in this account in the current region. |
| Rate of CreateMlflowApp requests | Each supported Region: 1 | No | The maximum number of CreateMlflowApp requests that you can make, per second, in this account in the current region. |
| Rate of CreateModel requests | Each supported Region: 1 | No | The maximum number of CreateModel requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateNotebookInstance requests | Each supported Region: 1 | No | The maximum number of CreateNotebookInstance requests that you can make, per second, in this account in the current region. |
| Rate of CreateNotebookInstanceLifecycleConfig requests | Each supported Region: 1 | No | The maximum number of CreateNotebookInstanceLifecycleConfig requests that you can make, per second, in this account in the current region. |
| Rate of CreatePresignedMlflowAppUrl requests | Each supported Region: 5 | No | The maximum number of CreatePresignedMlflowAppUrl requests that you can make, per second, in this account in the current region. |
| Rate of CreatePresignedNotebookInstanceUrl requests | Each supported Region: 1 | No | The maximum number of CreatePresignedNotebookInstanceUrl requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of CreateProcessingJob requests | Each supported Region: 1 | No | The maximum number of CreateProcessingJob requests that you can make, per second, in this account in the current region. |
| Rate of CreateStudioLifecycleConfig requests | Each supported Region: 1 | No | The maximum number of CreateStudioLifecycleConfig requests that you can make, per second, in this account in the current region. |
| Rate of CreateTrainingJob requests | Each supported Region: 1 | No | The maximum number of CreateTrainingJob requests that you can make, per second, in this account in the current region. |
| Rate of CreateTransformJob requests | Each supported Region: 1 | No | The maximum number of CreateTransformJob requests that you can make, per second, in this account in the current region. |
| Rate of CreateWorkforce requests | Each supported Region: 1 | No | The maximum number of CreateWorkforce requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of CreateWorkteam requests | Each supported Region: 1 | No | The maximum number of CreateWorkteam requests that you can make, per second, in this account in the current region. |
| Rate of DeleteEndpoint requests | Each supported Region: 1 | No | The maximum number of DeleteEndpoint requests that you can make, per second, in this account in the current region. |
| Rate of DeleteEndpointConfig requests | Each supported Region: 1 | No | The maximum number of DeleteEndpointConfig requests that you can make, per second, in this account in the current region. |
| Rate of DeleteFlowDefinition requests | Each supported Region: 1 | No | The maximum number of DeleteFlowDefinition requests that you can make, per second, in this account in the current region. |
| Rate of DeleteHumanTaskUi requests | Each supported Region: 1 | No | The maximum number of DeleteHumanTaskUi requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteMlflowApp requests | Each supported Region: 2 | No | The maximum number of DeleteMlflowApp requests that you can make, per second, in this account in the current region. |
| Rate of DeleteModel requests | Each supported Region: 1 | No | The maximum number of DeleteModel requests that you can make, per second, in this account in the current region. |
| Rate of DeleteNotebookInstance requests | Each supported Region: 1 | No | The maximum number of DeleteNotebookInstance requests that you can make, per second, in this account in the current region. |
| Rate of DeleteNotebookInstanceLifecycleConfig requests | Each supported Region: 1 | No | The maximum number of DeleteNotebookInstanceLifecycleConfig requests that you can make, per second, in this account in the current region. |
| Rate of DeleteStudioLifecycleConfig requests | Each supported Region: 1 | No | The maximum number of DeleteStudioLifecycleConfig requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DeleteTrainingJob requests | Each supported Region: 1 | No | The maximum number of DeleteTrainingJob requests that you can make, per second, in this account in the current region. |
| Rate of DeleteTransformJob requests | Each supported Region: 1 | No | The maximum number of DeleteTransformJob requests that you can make, per second, in this account in the current region. |
| Rate of DeleteWorkforce requests | Each supported Region: 1 | No | The maximum number of DeleteWorkforce requests that you can make, per second, in this account in the current region. |
| Rate of DeleteWorkteam requests | Each supported Region: 1 | No | The maximum number of DeleteWorkteam requests that you can make, per second, in this account in the current region. |
| Rate of DescribeEndpoint requests | Each supported Region: 5 | No | The maximum number of DescribeEndpoint requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeEndpointConfig requests | Each supported Region: 5 | No | The maximum number of DescribeEndpointConfig requests that you can make, per second, in this account in the current region. |
| Rate of DescribeFlowDefinition requests | Each supported Region: 5 | No | The maximum number of DescribeFlowDefinition requests that you can make, per second, in this account in the current region. |
| Rate of DescribeHumanTaskUi requests | Each supported Region: 5 | No | The maximum number of DescribeHumanTaskUi requests that you can make, per second, in this account in the current region. |
| Rate of DescribeLabelingJob requests | Each supported Region: 5 | No | The maximum number of DescribeLabelingJob requests that you can make, per second, in this account in the current region. |
| Rate of DescribeMlflowApp requests | Each supported Region: 1 | No | The maximum number of DescribeMlflowApp requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeModel requests | Each supported Region: 5 | No | The maximum number of DescribeModel requests that you can make, per second, in this account in the current region. |
| Rate of DescribeNotebookInstanceLifecycleConfig requests | Each supported Region: 5 | No | The maximum number of DescribeNotebookInstanceLifecycleConfig requests that you can make, per second, in this account in the current region. |
| Rate of DescribeStudioLifecycleConfig requests | Each supported Region: 5 | No | The maximum number of DescribeStudioLifecycleConfig requests that you can make, per second, in this account in the current region. |
| Rate of DescribeSubscribedWorkteam requests | Each supported Region: 1 | No | The maximum number of DescribeSubscribedWorkteam requests that you can make, per second, in this account in the current region. |
| Rate of DescribeTrainingJob requests | Each supported Region: 5 | No | The maximum number of DescribeTrainingJob requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of DescribeTransformJob requests | Each supported Region: 5 | No | The maximum number of DescribeTransformJob requests that you can make, per second, in this account in the current region. |
| Rate of DescribeWorkforce requests | Each supported Region: 5 | No | The maximum number of DescribeWorkforce requests that you can make, per second, in this account in the current region. |
| Rate of DescribeWorkteam requests | Each supported Region: 5 | No | The maximum number of DescribeWorkteam requests that you can make, per second, in this account in the current region. |
| Rate of InvokeEndpoint requests | Each supported Region: 10,000 | No | The maximum number of InvokeEndpoint requests that you can make, per second, in this account in the current region. |
| Rate of ListEndpointConfigs requests | Each supported Region: 2 | No | The maximum number of ListEndpointConfigs requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListEndpoints requests | Each supported Region: 2 | No | The maximum number of ListEndpoints requests that you can make, per second, in this account in the current region. |
| Rate of ListFlowDefinitions requests | Each supported Region: 2 | No | The maximum number of ListFlowDefinitions requests that you can make, per second, in this account in the current region. |
| Rate of ListHumanTaskUis requests | Each supported Region: 2 | No | The maximum number of ListHumanTaskUis requests that you can make, per second, in this account in the current region. |
| Rate of ListLabelingJobs requests | Each supported Region: 2 | No | The maximum number of ListLabelingJobs requests that you can make, per second, in this account in the current region. |
| Rate of ListLabelingJobsForWorkteam requests | Each supported Region: 2 | No | The maximum number of ListLabelingJobsForWorkteam requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListMlflowApps requests | Each supported Region: 1 | No | The maximum number of ListMlflowApps requests that you can make, per second, in this account in the current region. |
| Rate of ListModels requests | Each supported Region: 2 | No | The maximum number of ListModels requests that you can make, per second, in this account in the current region. |
| Rate of ListMonitoringAlertHistory requests | Each supported Region: 2 | No | The maximum number of ListMonitoringAlertHistory requests that you can make, per second, in this account in the current region. |
| Rate of ListMonitoringAlerts requests | Each supported Region: 2 | No | The maximum number of ListMonitoringAlerts requests that you can make, per second, in this account in the current region. |
| Rate of ListNotebookInstanceLifecycleConfigs requests | Each supported Region: 2 | No | The maximum number of ListNotebookInstanceLifecycleConfigs requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListNotebookInstances requests | Each supported Region: 2 | No | The maximum number of ListNotebookInstances requests that you can make, per second, in this account in the current region. |
| Rate of ListStudioLifecycleConfigs requests | Each supported Region: 2 | No | The maximum number of ListStudioLifecycleConfigs requests that you can make, per second, in this account in the current region. |
| Rate of ListSubscribedWorkteams requests | Each supported Region: 2 | No | The maximum number of ListSubscribedWorkteams requests that you can make, per second, in this account in the current region. |
| Rate of ListTrainingJobs requests | Each supported Region: 2 | No | The maximum number of ListTrainingJobs requests that you can make, per second, in this account in the current region. |
| Rate of ListTransformJobs requests | Each supported Region: 2 | No | The maximum number of ListTransformJobs requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of ListWorkforces requests | Each supported Region: 2 | No | The maximum number of ListWorkforces requests that you can make, per second, in this account in the current region. |
| Rate of ListWorkteams requests | Each supported Region: 2 | No | The maximum number of ListWorkteams requests that you can make, per second, in this account in the current region. |
| Rate of StartNotebookInstance requests | Each supported Region: 1 | No | The maximum number of StartNotebookInstance requests that you can make, per second, in this account in the current region. |
| Rate of StopLabelingJob requests | Each supported Region: 1 | No | The maximum number of StopLabelingJob requests that you can make, per second, in this account in the current region. |
| Rate of StopNotebookInstance requests | Each supported Region: 1 | No | The maximum number of StopNotebookInstance requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StopTrainingJob requests | Each supported Region: 1 | No | The maximum number of StopTrainingJob requests that you can make, per second, in this account in the current region. |
| Rate of StopTransformJob requests | Each supported Region: 1 | No | The maximum number of StopTransformJob requests that you can make, per second, in this account in the current region. |
| Rate of UpdateEndpoint requests | Each supported Region: 1 | No | The maximum number of UpdateEndpoint requests that you can make, per second, in this account in the current region. |
| Rate of UpdateEndpointWeightsAndCapacities requests | Each supported Region: 1 | No | The maximum number of UpdateEndpointWeightsAndCapacities requests that you can make, per second, in this account in the current region. |
| Rate of UpdateMlflowApp requests | Each supported Region: 1 | No | The maximum number of UpdateMlflowApp requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of UpdateMonitoringAlert requests | Each supported Region: 1 | No | The maximum number of UpdateMonitoringAlert requests that you can make, per second, in this account in the current region. |
| Rate of UpdateNotebookInstance requests | Each supported Region: 1 | No | The maximum number of UpdateNotebookInstance requests that you can make, per second, in this account in the current region. |
| Rate of UpdateNotebookInstanceLifecycleConfig requests | Each supported Region: 1 | No | The maximum number of UpdateNotebookInstanceLifecycleConfig requests that you can make, per second, in this account in the current region. |
| Rate of UpdateTrainingJob requests | Each supported Region: 1 | No | The maximum number of UpdateTrainingJob requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of UpdateWorkforce requests | Each supported Region: 1 | No | The maximum number of UpdateWorkforce requests that you can make, per second, in this account in the current region. |
| Rate of UpdateWorkteam requests | Each supported Region: 1 | No | The maximum number of UpdateWorkteam requests that you can make, per second, in this account in the current region. |
| SageMaker Profiler Apps running on system instances | Each supported Region: 5 | Yes | SageMaker Profiler Apps running on system instances |
| Size of EBS volume for a processing job instance | Each supported Region: 1,024 | Yes | Size of EBS volume for a processing job instance |
| Size of EBS volume for a training job instance | Each supported Region: 1,024 | Yes | Size of EBS volume for a training job instance |
| Size of EBS volume for a transform job instance | Each supported Region: 1,024 | No | Size of EBS volume for a transform job instance |
| Small-sized MLflow Tracking Server usage | Each supported Region: 5 | Yes | Small-sized MLflow Tracking Server usage |
| Studio CodeEditor Apps running on ml.c5.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c5.12x large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.c5.18xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c5.18x large instances |
| Studio CodeEditor Apps running on ml.c5.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c5.24x large instances |
| Studio CodeEditor Apps running on ml.c5.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c5.2xlarge instances |
| Studio CodeEditor Apps running on ml.c5.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c5.4xlarge instances |
| Studio CodeEditor Apps running on ml.c5.9xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c5.9xlarge instances |
| Studio CodeEditor Apps running on ml.c5.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c5.large instances |
| Studio CodeEditor Apps running on ml.c5.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c5.xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6i.12 xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6i.16 xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.c6i.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6i.24xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c6i.2xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.32xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6i.32xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6i.4xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6i.8xlarge instances |
| Studio CodeEditor Apps running on ml.c6i.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c6i.large instances |
| Studio CodeEditor Apps running on ml.c6i.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c6i.xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6id.12xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6id.16xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.c6id.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6id.24xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c6id.2xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.32xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6id.32xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6id.4xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c6id.8xlarge instances |
| Studio CodeEditor Apps running on ml.c6id.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c6id.large instances |
| Studio CodeEditor Apps running on ml.c6id.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c6id.xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c7i.12xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c7i.16xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.c7i.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c7i.24xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c7i.2xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c7i.48xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c7i.4xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.c7i.8xlarge instances |
| Studio CodeEditor Apps running on ml.c7i.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c7i.large instances |
| Studio CodeEditor Apps running on ml.c7i.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.c7i.xlarge instances |
| Studio CodeEditor Apps running on ml.g4dn.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g4dn.12xlarge instances |
| Studio CodeEditor Apps running on ml.g4dn.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g4dn.16xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.g4dn.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g4dn.2xlarge instances |
| Studio CodeEditor Apps running on ml.g4dn.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g4dn.4xlarge instances |
| Studio CodeEditor Apps running on ml.g4dn.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g4dn.8xlarge instances |
| Studio CodeEditor Apps running on ml.g4dn.xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g4dn.xlarge instances |
| Studio CodeEditor Apps running on ml.g5.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.12xlarge instances |
| Studio CodeEditor Apps running on ml.g5.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.16xlarge instances |
| Studio CodeEditor Apps running on ml.g5.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.24xlarge instances |
| Studio CodeEditor Apps running on ml.g5.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.2xlarge instances |
| Studio CodeEditor Apps running on ml.g5.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.48xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.g5.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.4xlarge instances |
| Studio CodeEditor Apps running on ml.g5.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.8xlarge instances |
| Studio CodeEditor Apps running on ml.g5.xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g5.xlarge instances |
| Studio CodeEditor Apps running on ml.g6.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.12x large instances |
| Studio CodeEditor Apps running on ml.g6.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.16x large instances |
| Studio CodeEditor Apps running on ml.g6.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.24x large instances |
| Studio CodeEditor Apps running on ml.g6.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.2xlarge instances |
| Studio CodeEditor Apps running on ml.g6.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.48x large instances |
| Studio CodeEditor Apps running on ml.g6.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.4xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.g6.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.8xlarge instances |
| Studio CodeEditor Apps running on ml.g6.xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6.xlarge instances |
| Studio CodeEditor Apps running on ml.g6e.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.12 xlarge instances |
| Studio CodeEditor Apps running on ml.g6e.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.16 xlarge instances |
| Studio CodeEditor Apps running on ml.g6e.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.24 xlarge instances |
| Studio CodeEditor Apps running on ml.g6e.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.2x large instances |
| Studio CodeEditor Apps running on ml.g6e.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.48 xlarge instances |
| Studio CodeEditor Apps running on ml.g6e.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.4x large instances |
| Studio CodeEditor Apps running on ml.g6e.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.8x large instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.g6e.xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g6e.xlarge instances |
| Studio CodeEditor Apps running on ml.g7.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7.12x large instances |
| Studio CodeEditor Apps running on ml.g7.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7.24x large instances |
| Studio CodeEditor Apps running on ml.g7.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7.2xlarge instances |
| Studio CodeEditor Apps running on ml.g7.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7.48x large instances |
| Studio CodeEditor Apps running on ml.g7.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7.4xlarge instances |
| Studio CodeEditor Apps running on ml.g7.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7.8xlarge instances |
| Studio CodeEditor Apps running on ml.g7e.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7e.12 xlarge instances |
| Studio CodeEditor Apps running on ml.g7e.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7e.24 xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.g7e.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7e.2xlarge instances |
| Studio CodeEditor Apps running on ml.g7e.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7e.48xlarge instances |
| Studio CodeEditor Apps running on ml.g7e.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7e.4xlarge instances |
| Studio CodeEditor Apps running on ml.g7e.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.g7e.8xlarge instances |
| Studio CodeEditor Apps running on ml.m5.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5.12xlarge instances |
| Studio CodeEditor Apps running on ml.m5.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5.16xlarge instances |
| Studio CodeEditor Apps running on ml.m5.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5.24xlarge instances |
| Studio CodeEditor Apps running on ml.m5.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m5.2xlarge instances |
| Studio CodeEditor Apps running on ml.m5.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5.4xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.m5.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5.8xlarge instances |
| Studio CodeEditor Apps running on ml.m5.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m5.large instances |
| Studio CodeEditor Apps running on ml.m5.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m5.xlarge instances |
| Studio CodeEditor Apps running on ml.m5d.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5d.12xlarge instances |
| Studio CodeEditor Apps running on ml.m5d.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5d.16xlarge instances |
| Studio CodeEditor Apps running on ml.m5d.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5d.24xlarge instances |
| Studio CodeEditor Apps running on ml.m5d.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5d.2xlarge instances |
| Studio CodeEditor Apps running on ml.m5d.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5d.4xlarge instances |
| Studio CodeEditor Apps running on ml.m5d.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m5d.8xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.m5d.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m5d.large instances |
| Studio CodeEditor Apps running on ml.m5d.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m5d.xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6i.12xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6i.16xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6i.24xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m6i.2xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.32xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6i.32xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6i.4xlarge instances |
| Studio CodeEditor Apps running on ml.m6i.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6i.8xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.m6i.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m6i.large instances |
| Studio CodeEditor Apps running on ml.m6i.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m6i.xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6id.12xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6id.16xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6id.24xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m6id.2xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.32xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6id.32xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6id.4xlarge instances |
| Studio CodeEditor Apps running on ml.m6id.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m6id.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.m6id.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m6id.large instances |
| Studio CodeEditor Apps running on ml.m6id.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m6id.xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m7i.12xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m7i.16xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m7i.24xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.2xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m7i.2xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m7i.48xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m7i.4xlarge instances |
| Studio CodeEditor Apps running on ml.m7i.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.m7i.8xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.m7i.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m7i.large instances |
| Studio CodeEditor Apps running on ml.m7i.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.m7i.xlarge instances |
| Studio CodeEditor Apps running on ml.p3.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p3.16x large instances |
| Studio CodeEditor Apps running on ml.p3.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p3.2xlarge instances |
| Studio CodeEditor Apps running on ml.p3.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p3.8xlarge instances |
| Studio CodeEditor Apps running on ml.p3dn.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p3dn.2 4xlarge instances |
| Studio CodeEditor Apps running on ml.p4d.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p4d.24 xlarge instances |
| Studio CodeEditor Apps running on ml.p4de.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p4de.2 4xlarge instances |
| Studio CodeEditor Apps running on ml.p5.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p5.48x large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.p5.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p5.4xlarge instances |
| Studio CodeEditor Apps running on ml.p5en.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p5en.48xlarge instances |
| Studio CodeEditor Apps running on ml.p6-b200.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.p6-b200.48xlarge instances |
| Studio CodeEditor Apps running on ml.r5.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r5.12xlarge instances |
| Studio CodeEditor Apps running on ml.r5.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r5.16xlarge instances |
| Studio CodeEditor Apps running on ml.r5.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r5.24xlarge instances |
| Studio CodeEditor Apps running on ml.r5.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r5.2xlarge instances |
| Studio CodeEditor Apps running on ml.r5.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r5.4xlarge instances |
| Studio CodeEditor Apps running on ml.r5.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r5.8xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.r5.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r5.large instances |
| Studio CodeEditor Apps running on ml.r5.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r5.xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.12 xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.16 xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.24 xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.2xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.32xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.32 xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.4xlarge instances |
| Studio CodeEditor Apps running on ml.r6i.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6i.8xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.r6i.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r6i.large instances |
| Studio CodeEditor Apps running on ml.r6i.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r6i.xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.12xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.16xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.24xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.2xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.32xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.32xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.4xlarge instances |
| Studio CodeEditor Apps running on ml.r6id.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r6id.8xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio CodeEditor Apps running on ml.r6id.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r6id.large instances |
| Studio CodeEditor Apps running on ml.r6id.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r6id.xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.12xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.12 xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.16xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.16 xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.24xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.24 xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.2xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.2xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.48xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.48 xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.4xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.4xlarge instances |
| Studio CodeEditor Apps running on ml.r7i.8xlarge instances | Each supported Region: 0 | Yes | Studio CodeEditor Apps running on ml.r7i.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio CodeEditor Apps running on ml.r7i.large instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r7i.large instances |
| Studio CodeEditor Apps running on ml.r7i.xlarge instances | Each supported Region: 1 | Yes | Studio CodeEditor Apps running on ml.r7i.xlarge instances |
| Studio CodeEditor Apps running on ml.t3.2xlarge instances | Each supported Region: 6 | Yes | Studio CodeEditor Apps running on ml.t3.2xlarge instances |
| Studio CodeEditor Apps running on ml.t3.large instances | Each supported Region: 6 | Yes | Studio CodeEditor Apps running on ml.t3.large instances |
| Studio CodeEditor Apps running on ml.t3.medium instances | Each supported Region: 6 | Yes | Studio CodeEditor Apps running on ml.t3.medium instances |
| Studio CodeEditor Apps running on ml.t3.xlarge instances | Each supported Region: 6 | Yes | Studio CodeEditor Apps running on ml.t3.xlarge instances |
| Studio Jupyter Apps running on system instances | Each supported Region: 10 | Yes | Studio Jupyter Apps running on system instances |
| Studio JupyterLab Apps running on ml.c5.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c5.12xlarge instances |
| Studio JupyterLab Apps running on ml.c5.18xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c5.18xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.c5.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c5.24xlarge instances |
| Studio JupyterLab Apps running on ml.c5.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c5.2xlarge instances |
| Studio JupyterLab Apps running on ml.c5.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c5.4xlarge instances |
| Studio JupyterLab Apps running on ml.c5.9xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c5.9xlarge instances |
| Studio JupyterLab Apps running on ml.c5.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c5.large instances |
| Studio JupyterLab Apps running on ml.c5.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c5.xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6i.12xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6i.16xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6i.24xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.c6i.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c6i.2xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.32xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6i.32xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6i.4xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6i.8xlarge instances |
| Studio JupyterLab Apps running on ml.c6i.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c6i.large instances |
| Studio JupyterLab Apps running on ml.c6i.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c6i.xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6id.12xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6id.16xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6id.24xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.c6id.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c6id.2xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.32xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6id.32xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6id.4xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c6id.8xlarge instances |
| Studio JupyterLab Apps running on ml.c6id.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c6id.large instances |
| Studio JupyterLab Apps running on ml.c6id.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c6id.xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c7i.12xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c7i.16xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c7i.24xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.c7i.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c7i.2xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c7i.48xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c7i.4xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.c7i.8xlarge instances |
| Studio JupyterLab Apps running on ml.c7i.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c7i.large instances |
| Studio JupyterLab Apps running on ml.c7i.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.c7i.xlarge instances |
| Studio JupyterLab Apps running on ml.g4dn.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g4dn.12xlarge instances |
| Studio JupyterLab Apps running on ml.g4dn.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g4dn.16xlarge instances |
| Studio JupyterLab Apps running on ml.g4dn.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g4dn.2xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.g4dn.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g4dn.4xlarge instances |
| Studio JupyterLab Apps running on ml.g4dn.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g4dn.8xlarge instances |
| Studio JupyterLab Apps running on ml.g4dn.xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g4dn.xlarge instances |
| Studio JupyterLab Apps running on ml.g5.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.12xlarge instances |
| Studio JupyterLab Apps running on ml.g5.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.16xlarge instances |
| Studio JupyterLab Apps running on ml.g5.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.24xlarge instances |
| Studio JupyterLab Apps running on ml.g5.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.2xlarge instances |
| Studio JupyterLab Apps running on ml.g5.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.48xlarge instances |
| Studio JupyterLab Apps running on ml.g5.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.4xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.g5.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.8xlarge instances |
| Studio JupyterLab Apps running on ml.g5.xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g5.xlarge instances |
| Studio JupyterLab Apps running on ml.g6.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.12x large instances |
| Studio JupyterLab Apps running on ml.g6.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.16x large instances |
| Studio JupyterLab Apps running on ml.g6.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.24x large instances |
| Studio JupyterLab Apps running on ml.g6.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.2xlarge instances |
| Studio JupyterLab Apps running on ml.g6.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.48x large instances |
| Studio JupyterLab Apps running on ml.g6.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.4xlarge instances |
| Studio JupyterLab Apps running on ml.g6.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.g6.xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6.xlarge instances |
| Studio JupyterLab Apps running on ml.g6e.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.12 xlarge instances |
| Studio JupyterLab Apps running on ml.g6e.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.16 xlarge instances |
| Studio JupyterLab Apps running on ml.g6e.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.24 xlarge instances |
| Studio JupyterLab Apps running on ml.g6e.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.2x large instances |
| Studio JupyterLab Apps running on ml.g6e.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.48 xlarge instances |
| Studio JupyterLab Apps running on ml.g6e.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.4x large instances |
| Studio JupyterLab Apps running on ml.g6e.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.8x large instances |
| Studio JupyterLab Apps running on ml.g6e.xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g6e.xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.g7.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7.12xlarge instances |
| Studio JupyterLab Apps running on ml.g7.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7.24xlarge instances |
| Studio JupyterLab Apps running on ml.g7.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7.2xlarge instances |
| Studio JupyterLab Apps running on ml.g7.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7.48xlarge instances |
| Studio JupyterLab Apps running on ml.g7.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7.4xlarge instances |
| Studio JupyterLab Apps running on ml.g7.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7.8xlarge instances |
| Studio JupyterLab Apps running on ml.g7e.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7e.12xlarge instances |
| Studio JupyterLab Apps running on ml.g7e.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7e.24xlarge instances |
| Studio JupyterLab Apps running on ml.g7e.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7e.2xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.g7e.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7e.48xlarge instances |
| Studio JupyterLab Apps running on ml.g7e.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7e.4xlarge instances |
| Studio JupyterLab Apps running on ml.g7e.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.g7e.8xlarge instances |
| Studio JupyterLab Apps running on ml.m5.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5.12xlarge instances |
| Studio JupyterLab Apps running on ml.m5.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5.16xlarge instances |
| Studio JupyterLab Apps running on ml.m5.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5.24xlarge instances |
| Studio JupyterLab Apps running on ml.m5.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m5.2xlarge instances |
| Studio JupyterLab Apps running on ml.m5.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5.4xlarge instances |
| Studio JupyterLab Apps running on ml.m5.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5.8xlarge instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.m5.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m5.large instances |
| Studio JupyterLab Apps running on ml.m5.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m5.xlarge instances |
| Studio JupyterLab Apps running on ml.m5d.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5d.12 xlarge instances |
| Studio JupyterLab Apps running on ml.m5d.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5d.16 xlarge instances |
| Studio JupyterLab Apps running on ml.m5d.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5d.24 xlarge instances |
| Studio JupyterLab Apps running on ml.m5d.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5d.2x large instances |
| Studio JupyterLab Apps running on ml.m5d.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5d.4x large instances |
| Studio JupyterLab Apps running on ml.m5d.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m5d.8x large instances |
| Studio JupyterLab Apps running on ml.m5d.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m5d.large instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.m5d.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m5d.xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6i.12xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6i.16xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6i.24xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m6i.2xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.32xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6i.32xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6i.4xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6i.8xlarge instances |
| Studio JupyterLab Apps running on ml.m6i.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m6i.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.m6i.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m6i.xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6id.12xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6id.16xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6id.24xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m6id.2xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.32xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6id.32xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6id.4xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m6id.8xlarge instances |
| Studio JupyterLab Apps running on ml.m6id.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m6id.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.m6id.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m6id.xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m7i.12xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m7i.16xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m7i.24xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.2xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m7i.2xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m7i.48xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m7i.4xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.m7i.8xlarge instances |
| Studio JupyterLab Apps running on ml.m7i.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m7i.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.m7i.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.m7i.xlarge instances |
| Studio JupyterLab Apps running on ml.p3.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p3.16x large instances |
| Studio JupyterLab Apps running on ml.p3.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p3.2xlarge instances |
| Studio JupyterLab Apps running on ml.p3.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p3.8xlarge instances |
| Studio JupyterLab Apps running on ml.p3dn.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p3dn.2 4xlarge instances |
| Studio JupyterLab Apps running on ml.p4d.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p4d.24 xlarge instances |
| Studio JupyterLab Apps running on ml.p4de.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p4de.2 4xlarge instances |
| Studio JupyterLab Apps running on ml.p5.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p5.48x large instances |
| Studio JupyterLab Apps running on ml.p5.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p5.4xlarge instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.p5en.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p5en.48xlarge instances |
| Studio JupyterLab Apps running on ml.p6-b200.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.p6-b200.48xlarge instances |
| Studio JupyterLab Apps running on ml.r5.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r5.12xlarge instances |
| Studio JupyterLab Apps running on ml.r5.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r5.16xlarge instances |
| Studio JupyterLab Apps running on ml.r5.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r5.24xlarge instances |
| Studio JupyterLab Apps running on ml.r5.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r5.2xlarge instances |
| Studio JupyterLab Apps running on ml.r5.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r5.4xlarge instances |
| Studio JupyterLab Apps running on ml.r5.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r5.8xlarge instances |
| Studio JupyterLab Apps running on ml.r5.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r5.large instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.r5.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r5.xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.12xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.16xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.24xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.2xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.32xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.32xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.4xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6i.8xlarge instances |
| Studio JupyterLab Apps running on ml.r6i.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r6i.large instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio JupyterLab Apps running on ml.r6i.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r6i.xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.12xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.16xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.24xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.2xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.32xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.32xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.4xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r6id.8xlarge instances |
| Studio JupyterLab Apps running on ml.r6id.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r6id.large instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.r6id.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r6id.xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.12xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.12 xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.16xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.16 xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.24xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.24 xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.2xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.2xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.48xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.48 xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.4xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.4xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.8xlarge instances | Each supported Region: 0 | Yes | Studio JupyterLab Apps running on ml.r7i.8xlarge instances |
| Studio JupyterLab Apps running on ml.r7i.large instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r7i.large instances |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio JupyterLab Apps running on ml.r7i.xlarge instances | Each supported Region: 1 | Yes | Studio JupyterLab Apps running on ml.r7i.xlarge instances |
| Studio JupyterLab Apps running on ml.t3.2xlarge instances | Each supported Region: 6 | Yes | Studio JupyterLab Apps running on ml.t3.2xlarge instances |
| Studio JupyterLab Apps running on ml.t3.large instances | Each supported Region: 6 | Yes | Studio JupyterLab Apps running on ml.t3.large instances |
| Studio JupyterLab Apps running on ml.t3.medium instances | Each supported Region: 6 | Yes | Studio JupyterLab Apps running on ml.t3.medium instances |
| Studio JupyterLab Apps running on ml.t3.xlarge instances | Each supported Region: 6 | Yes | Studio JupyterLab Apps running on ml.t3.xlarge instances |
| Studio KernelGateway Apps running on ml.c5.12xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.c5.12xlarge instance |
| Studio KernelGateway Apps running on ml.c5.18xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.c5.18xlarge instance |
| Studio KernelGateway Apps running on ml.c5.24xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.c5.24xlarge instance |
| Studio KernelGateway Apps running on ml.c5.2xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.c5.2xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio KernelGateway Apps running on ml.c5.4xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.c5.4xlarge instance |
| Studio KernelGateway Apps running on ml.c5.9xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.c5.9xlarge instance |
| Studio KernelGateway Apps running on ml.c5.large instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.c5.large instance |
| Studio KernelGateway Apps running on ml.c5.xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.c5.xlarge instance |
| Studio KernelGateway Apps running on ml.g4dn.12xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g4dn.12xlarge instance |
| Studio KernelGateway Apps running on ml.g4dn.16xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g4dn.16xlarge instance |
| Studio KernelGateway Apps running on ml.g4dn.2xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g4dn.2xlarge instance |
| Studio KernelGateway Apps running on ml.g4dn.4xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g4dn.4xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio KernelGateway Apps running on ml.g4dn.8xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g4dn.8xlarge instance |
| Studio KernelGateway Apps running on ml.g4dn.xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g4dn.xlarge instance |
| Studio KernelGateway Apps running on ml.g5.12xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.12xlarge instance |
| Studio KernelGateway Apps running on ml.g5.16xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.16xlarge instance |
| Studio KernelGateway Apps running on ml.g5.24xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.24xlarge instance |
| Studio KernelGateway Apps running on ml.g5.2xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.2xlarge instance |
| Studio KernelGateway Apps running on ml.g5.48xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.48xlarge instance |
| Studio KernelGateway Apps running on ml.g5.4xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.4xlarge instance |
| Studio KernelGateway Apps running on ml.g5.8xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.8xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio KernelGateway Apps running on ml.g5.xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.g5.xlarge instance |
| Studio KernelGateway Apps running on ml.m5.12xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5.12xlarge instance |
| Studio KernelGateway Apps running on ml.m5.16xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5.16xlarge instance |
| Studio KernelGateway Apps running on ml.m5.24xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5.24xlarge instance |
| Studio KernelGateway Apps running on ml.m5.2xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5.2xlarge instance |
| Studio KernelGateway Apps running on ml.m5.4xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5.4xlarge instance |
| Studio KernelGateway Apps running on ml.m5.8xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5.8xlarge instance |
| Studio KernelGateway Apps running on ml.m5.large instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5.large instance |
| Studio KernelGateway Apps running on ml.m5.xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5.xlarge instance |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Studio KernelGateway Apps running on ml.m5d.12xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5d.12xlarge instance |
| Studio KernelGateway Apps running on ml.m5d.16xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5d.16xlarge instance |
| Studio KernelGateway Apps running on ml.m5d.24xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5d.24xlarge instance |
| Studio KernelGateway Apps running on ml.m5d.2xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5d.2xlarge instance |
| Studio KernelGateway Apps running on ml.m5d.4xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5d.4xlarge instance |
| Studio KernelGateway Apps running on ml.m5d.8xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5d.8xlarge instance |
| Studio KernelGateway Apps running on ml.m5d.large instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.m5d.large instance |
| Studio KernelGateway Apps running on ml.m5d.xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.m5d.xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio KernelGateway Apps running on ml.p3.16xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.p3.16xlarge instance |
| Studio KernelGateway Apps running on ml.p3.2xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.p3.2xlarge instance |
| Studio KernelGateway Apps running on ml.p3.8xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.p3.8xlarge instance |
| Studio KernelGateway Apps running on ml.p3dn.24xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.p3dn.24xlarge instance |
| Studio KernelGateway Apps running on ml.p4d.24xlarge instances | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.p4d.24xlarge instances |
| Studio KernelGateway Apps running on ml.p4de.24xlarge instances | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.p4de.24xlarge instances |
| Studio KernelGateway Apps running on ml.r5.12xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.r5.12xlarge instance |
| Studio KernelGateway Apps running on ml.r5.16xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.r5.16xlarge instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio KernelGateway Apps running on ml.r5.24xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.r5.24xlarge instance |
| Studio KernelGateway Apps running on ml.r5.2xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.r5.2xlarge instance |
| Studio KernelGateway Apps running on ml.r5.4xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.r5.4xlarge instance |
| Studio KernelGateway Apps running on ml.r5.8xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.r5.8xlarge instance |
| Studio KernelGateway Apps running on ml.r5.large instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.r5.large instance |
| Studio KernelGateway Apps running on ml.r5.xlarge instance | Each supported Region: 1 | Yes | Studio KernelGateway Apps running on ml.r5.xlarge instance |
| Studio KernelGateway Apps running on ml.t3.2xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.t3.2xlarge instance |
| Studio KernelGateway Apps running on ml.t3.large instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.t3.large instance |
| Studio KernelGateway Apps running on ml.t3.medium instance | Each supported Region: 2 | Yes | Studio KernelGateway Apps running on ml.t3.medium instance |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Studio KernelGateway Apps running on ml.t3.xlarge instance | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.t3.xlarge instance |
| Studio KernelGateway Apps running on ml.trn1.2xlarge instances | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.trn1.2xlarge instances |
| Studio KernelGateway Apps running on ml.trn1.32xlarge instances | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.trn1.32xlarge instances |
| Studio KernelGateway Apps running on ml.trn1n.32xlarge instances | Each supported Region: 0 | Yes | Studio KernelGateway Apps running on ml.trn1n.32xlarge instances |
| TensorBoard Apps running on system instances | Each supported Region: 10 | Yes | TensorBoard Apps running on system instances |
| The maximum number of MLflow Apps that can be created. | Each supported Region: 3 | Yes | The maximum number of MLflow Apps that can be created. |
| Time at which pipeline executions time out | Each supported Region: 672 | No | Time at which pipeline executions time out |
| Total EBS volume size in GB across all Studio Spaces | Each supported Region: 500 | Yes | Total EBS volume size in GB across all Studio Spaces |
| Total EBS volume size in GB across all notebook instances | Each supported Region: 102,400 | Yes | Total EBS volume size in GB across all notebook instances |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Total Monitoring Schedules | Each supported Region: 10 | No | Total Monitoring Schedules |
| Total domains | Each supported Region: 2 | Yes | Total domains |
| Total number of Spot instances allowed across SageMaker HyperPod clusters | Each supported Region: 0 | Yes | Total number of Spot instances allowed across SageMaker HyperPod clusters |
| Total number of experiments allowed, excluding those automatically created by SageMaker | Each supported Region: 5,000 | Yes | Total number of experiments allowed, excluding those automatically created by SageMaker |
| Total number of instances allowed across SageMaker HyperPod clusters | Each supported Region: 0 | Yes | Total number of instances allowed across SageMaker HyperPod clusters |
| Total number of notebook instances | Each supported Region: 8 | Yes | Total number of notebook instances |
| Total number of trial components allowed from a SageMaker context, excluding those automatically created by SageMaker | Each supported Region: 20,000 | Yes | Total number of trial components allowed from a SageMaker context, excluding those automatically created by SageMaker |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Total number of trial components allowed in a single trial, excluding those automatically created by SageMaker | Each supported Region: 50 | Yes | Total number of trial components allowed in a single trial, excluding those automatically created by SageMaker |
| Total number of trials a single trial component can be associated to | Each supported Region: 500 | Yes | Total number of trials a single trial component can be associated to |
| Total number of trials allowed in a single experiment, excluding those automatically created by SageMaker | Each supported Region: 300 | Yes | Total number of trials allowed in a single experiment, excluding those automatically created by SageMaker |
| cluster-ml-r7i-12xlarge for cluster usage | Each supported Region: 0 | Yes | cluster-ml-r7i-12xlarge for cluster usage |
| ml.c4.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c4.2xlarge for endpoint usage |
| ml.c4.2xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c4.2xlarge for notebook instance usage |
| ml.c4.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c4.2xlarge for processing job usage |
| ml.c4.2xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.c4.2xlarge for spot training job usage |
| ml.c4.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.c4.2xlarge for training job usage |
| ml.c4.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c4.2xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c4.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c4.2xlarge for transform job usage |
| ml.c4.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for endpoint usage |
| ml.c4.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for notebook instance usage |
| ml.c4.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for processing job usage |
| ml.c4.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for spot training job usage |
| ml.c4.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for training job usage |
| ml.c4.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for training warm pool usage |
| ml.c4.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c4.4xlarge for transform job usage |
| ml.c4.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for endpoint usage |
| ml.c4.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for notebook instance usage |
| ml.c4.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for processing job usage |
| ml.c4.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c4.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for training job usage |
| ml.c4.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for training warm pool usage |
| ml.c4.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c4.8xlarge for transform job usage |
| ml.c4.large for endpoint usage | Each supported Region: 0 | Yes | ml.c4.large for endpoint usage |
| ml.c4.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c4.xlarge for endpoint usage |
| ml.c4.xlarge for notebook instance usage | Each supported Region: 2 | Yes | ml.c4.xlarge for notebook instance usage |
| ml.c4.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c4.xlarge for processing job usage |
| ml.c4.xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.c4.xlarge for spot training job usage |
| ml.c4.xlarge for training job usage | Each supported Region: 0 | Yes | ml.c4.xlarge for training job usage |
| ml.c4.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c4.xlarge for training warm pool usage |
| ml.c4.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c4.xlarge for transform job usage |
| ml.c5.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.12xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.12xlarge for cluster usage |
| ml.c5.18xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for cluster spot instance usage |
| ml.c5.18xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for cluster usage |
| ml.c5.18xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for endpoint usage |
| ml.c5.18xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for notebook instance usage |
| ml.c5.18xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for processing job usage |
| ml.c5.18xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for spot training job usage |
| ml.c5.18xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for training job usage |
| ml.c5.18xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for training warm pool usage |
| ml.c5.18xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c5.18xlarge for transform job usage |
| ml.c5.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.24xlarge for cluster spot instance usage |
| ml.c5.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.24xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for cluster spot instance usage |
| ml.c5.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for cluster usage |
| ml.c5.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for endpoint usage |
| ml.c5.2xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c5.2xlarge for notebook instance usage |
| ml.c5.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for processing job usage |
| ml.c5.2xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.c5.2xlarge for spot training job usage |
| ml.c5.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for training job usage |
| ml.c5.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for training warm pool usage |
| ml.c5.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c5.2xlarge for transform job usage |
| ml.c5.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for cluster spot instance usage |
| ml.c5.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for cluster usage |
| ml.c5.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for notebook instance usage |
| ml.c5.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for processing job usage |
| ml.c5.4xlarge for spot training job usage | Each supported Region: 1 | Yes | ml.c5.4xlarge for spot training job usage |
| ml.c5.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for training job usage |
| ml.c5.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for training warm pool usage |
| ml.c5.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c5.4xlarge for transform job usage |
| ml.c5.9xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for cluster spot instance usage |
| ml.c5.9xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for cluster usage |
| ml.c5.9xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for endpoint usage |
| ml.c5.9xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for notebook instance usage |
| ml.c5.9xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for processing job usage |
| ml.c5.9xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5.9xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for training job usage |
| ml.c5.9xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for training warm pool usage |
| ml.c5.9xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c5.9xlarge for transform job usage |
| ml.c5.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.large for cluster spot instance usage |
| ml.c5.large for cluster usage | Each supported Region: 0 | Yes | ml.c5.large for cluster usage |
| ml.c5.large for endpoint usage | Each supported Region: 0 | Yes | ml.c5.large for endpoint usage |
| ml.c5.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5.xlarge for cluster spot instance usage |
| ml.c5.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5.xlarge for cluster usage |
| ml.c5.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5.xlarge for endpoint usage |
| ml.c5.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c5.xlarge for notebook instance usage |
| ml.c5.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c5.xlarge for processing job usage |
| ml.c5.xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.c5.xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5.xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5.xlarge for training job usage |
| ml.c5.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5.xlarge for training warm pool usage |
| ml.c5.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c5.xlarge for transform job usage |
| ml.c5d.18xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5d.18xlarge for endpoint usage |
| ml.c5d.18xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c5d.18xlarge for notebook instance usage |
| ml.c5d.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5d.2xlarge for endpoint usage |
| ml.c5d.2xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c5d.2xlarge for notebook instance usage |
| ml.c5d.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5d.4xlarge for endpoint usage |
| ml.c5d.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c5d.4xlarge for notebook instance usage |
| ml.c5d.9xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5d.9xlarge for endpoint usage |
| ml.c5d.9xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c5d.9xlarge for notebook instance usage |
| ml.c5d.large for endpoint usage | Each supported Region: 0 | Yes | ml.c5d.large for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5d.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c5d.xlarge for endpoint usage |
| ml.c5d.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c5d.xlarge for notebook instance usage |
| ml.c5n.18xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5n.18xlarge for cluster spot instance usage |
| ml.c5n.18xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5n.18xlarge for cluster usage |
| ml.c5n.18xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5n.18xlarge for spot training job usage |
| ml.c5n.18xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5n.18xlarge for training job usage |
| ml.c5n.18xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5n.18xlarge for training warm pool usage |
| ml.c5n.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5n.2xlarge for cluster spot instance usage |
| ml.c5n.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5n.2xlarge for cluster usage |
| ml.c5n.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5n.2xlarge for spot training job usage |
| ml.c5n.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5n.2xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5n.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5n.2xlarge for training warm pool usage |
| ml.c5n.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5n.4xlarge for cluster spot instance usage |
| ml.c5n.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5n.4xlarge for cluster usage |
| ml.c5n.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5n.4xlarge for spot training job usage |
| ml.c5n.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5n.4xlarge for training job usage |
| ml.c5n.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5n.4xlarge for training warm pool usage |
| ml.c5n.9xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5n.9xlarge for cluster spot instance usage |
| ml.c5n.9xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c5n.9xlarge for cluster usage |
| ml.c5n.9xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5n.9xlarge for spot training job usage |
| ml.c5n.9xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5n.9xlarge for training job usage |
| ml.c5n.9xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5n.9xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c5n.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c5n.large for cluster spot instance usage |
| ml.c5n.large for cluster usage | Each supported Region: 0 | Yes | ml.c5n.large for cluster usage |
| ml.c5n.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c5n.xlarge for spot training job usage |
| ml.c5n.xlarge for training job usage | Each supported Region: 0 | Yes | ml.c5n.xlarge for training job usage |
| ml.c5n.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c5n.xlarge for training warm pool usage |
| ml.c6a.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.12xlarge for cluster spot instance usage |
| ml.c6a.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.12xlarge for cluster usage |
| ml.c6a.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.16xlarge for cluster spot instance usage |
| ml.c6a.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.16xlarge for cluster usage |
| ml.c6a.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.24xlarge for cluster spot instance usage |
| ml.c6a.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.24xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6a.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.2xlarge for cluster spot instance usage |
| ml.c6a.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.2xlarge for cluster usage |
| ml.c6a.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.32xlarge for cluster spot instance usage |
| ml.c6a.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.32xlarge for cluster usage |
| ml.c6a.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.48xlarge for cluster spot instance usage |
| ml.c6a.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.48xlarge for cluster usage |
| ml.c6a.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.4xlarge for cluster spot instance usage |
| ml.c6a.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.4xlarge for cluster usage |
| ml.c6a.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.8xlarge for cluster spot instance usage |
| ml.c6a.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.8xlarge for cluster usage |
| ml.c6a.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.large for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6a.large for cluster usage | Each supported Region: 0 | Yes | ml.c6a.large for cluster usage |
| ml.c6a.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6a.xlarge for cluster spot instance usage |
| ml.c6a.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6a.xlarge for cluster usage |
| ml.c6g.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.12xlarge for cluster spot instance usage |
| ml.c6g.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6g.12xlarge for cluster usage |
| ml.c6g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.12xlarge for endpoint usage |
| ml.c6g.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.16xlarge for cluster spot instance usage |
| ml.c6g.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6g.16xlarge for cluster usage |
| ml.c6g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.16xlarge for endpoint usage |
| ml.c6g.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.2xlarge for cluster spot instance usage |
| ml.c6g.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6g.2xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.2xlarge for endpoint usage |
| ml.c6g.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.4xlarge for cluster spot instance usage |
| ml.c6g.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6g.4xlarge for cluster usage |
| ml.c6g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.4xlarge for endpoint usage |
| ml.c6g.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.8xlarge for cluster spot instance usage |
| ml.c6g.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6g.8xlarge for cluster usage |
| ml.c6g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.8xlarge for endpoint usage |
| ml.c6g.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.large for cluster spot instance usage |
| ml.c6g.large for cluster usage | Each supported Region: 0 | Yes | ml.c6g.large for cluster usage |
| ml.c6g.large for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.large for endpoint usage |
| ml.c6g.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.medium for cluster spot instance usage |
| ml.c6g.medium for cluster usage | Each supported Region: 0 | Yes | ml.c6g.medium for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6g.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6g.xlarge for cluster spot instance usage |
| ml.c6g.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6g.xlarge for cluster usage |
| ml.c6g.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6g.xlarge for endpoint usage |
| ml.c6gd.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.12xlarge for endpoint usage |
| ml.c6gd.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.16xlarge for endpoint usage |
| ml.c6gd.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.2xlarge for endpoint usage |
| ml.c6gd.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.4xlarge for endpoint usage |
| ml.c6gd.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.8xlarge for endpoint usage |
| ml.c6gd.large for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.large for endpoint usage |
| ml.c6gd.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gd.xlarge for endpoint usage |
| ml.c6gn.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.12xlarge for endpoint usage |
| ml.c6gn.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.16xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6gn.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.2xlarge for endpoint usage |
| ml.c6gn.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.4xlarge for endpoint usage |
| ml.c6gn.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.8xlarge for endpoint usage |
| ml.c6gn.large for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.large for endpoint usage |
| ml.c6gn.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6gn.xlarge for endpoint usage |
| ml.c6i.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for cluster spot instance usage |
| ml.c6i.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for cluster usage |
| ml.c6i.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for endpoint usage |
| ml.c6i.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for notebook instance usage |
| ml.c6i.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for processing job usage |
| ml.c6i.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for spot training job usage |
| ml.c6i.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for training warm pool usage |
| ml.c6i.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.12xlarge for transform job usage |
| ml.c6i.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for cluster spot instance usage |
| ml.c6i.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for cluster usage |
| ml.c6i.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for endpoint usage |
| ml.c6i.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for notebook instance usage |
| ml.c6i.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for processing job usage |
| ml.c6i.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for spot training job usage |
| ml.c6i.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for training job usage |
| ml.c6i.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for training warm pool usage |
| ml.c6i.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.16xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for cluster spot instance usage |
| ml.c6i.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for cluster usage |
| ml.c6i.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for endpoint usage |
| ml.c6i.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for notebook instance usage |
| ml.c6i.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for processing job usage |
| ml.c6i.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for spot training job usage |
| ml.c6i.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for training job usage |
| ml.c6i.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for training warm pool usage |
| ml.c6i.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.24xlarge for transform job usage |
| ml.c6i.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for cluster spot instance usage |
| ml.c6i.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for endpoint usage |
| ml.c6i.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for notebook instance usage |
| ml.c6i.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for processing job usage |
| ml.c6i.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for spot training job usage |
| ml.c6i.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for training job usage |
| ml.c6i.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for training warm pool usage |
| ml.c6i.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.2xlarge for transform job usage |
| ml.c6i.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for cluster spot instance usage |
| ml.c6i.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for cluster usage |
| ml.c6i.32xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for endpoint usage |
| ml.c6i.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.32xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for processing job usage |
| ml.c6i.32xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for spot training job usage |
| ml.c6i.32xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for training job usage |
| ml.c6i.32xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for training warm pool usage |
| ml.c6i.32xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.32xlarge for transform job usage |
| ml.c6i.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for cluster spot instance usage |
| ml.c6i.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for cluster usage |
| ml.c6i.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for endpoint usage |
| ml.c6i.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for notebook instance usage |
| ml.c6i.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for processing job usage |
| ml.c6i.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for spot training job usage |
| ml.c6i.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for training warm pool usage |
| ml.c6i.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.4xlarge for transform job usage |
| ml.c6i.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for cluster spot instance usage |
| ml.c6i.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for cluster usage |
| ml.c6i.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for endpoint usage |
| ml.c6i.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for notebook instance usage |
| ml.c6i.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for processing job usage |
| ml.c6i.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for spot training job usage |
| ml.c6i.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for training job usage |
| ml.c6i.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for training warm pool usage |
| ml.c6i.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.8xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.large for cluster spot instance usage |
| ml.c6i.large for cluster usage | Each supported Region: 0 | Yes | ml.c6i.large for cluster usage |
| ml.c6i.large for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.large for endpoint usage |
| ml.c6i.large for notebook instance usage | Each supported Region: 1 | Yes | ml.c6i.large for notebook instance usage |
| ml.c6i.large for transform job usage | Each supported Region: 0 | Yes | ml.c6i.large for transform job usage |
| ml.c6i.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for cluster spot instance usage |
| ml.c6i.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for cluster usage |
| ml.c6i.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for endpoint usage |
| ml.c6i.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c6i.xlarge for notebook instance usage |
| ml.c6i.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for processing job usage |
| ml.c6i.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for spot training job usage |
| ml.c6i.xlarge for training job usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6i.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for training warm pool usage |
| ml.c6i.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c6i.xlarge for transform job usage |
| ml.c6id.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.12xlarge for notebook instance usage |
| ml.c6id.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.16xlarge for notebook instance usage |
| ml.c6id.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.24xlarge for notebook instance usage |
| ml.c6id.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.2xlarge for notebook instance usage |
| ml.c6id.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.32xlarge for notebook instance usage |
| ml.c6id.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.4xlarge for notebook instance usage |
| ml.c6id.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c6id.8xlarge for notebook instance usage |
| ml.c6id.large for notebook instance usage | Each supported Region: 1 | Yes | ml.c6id.large for notebook instance usage |
| ml.c6id.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c6id.xlarge for notebook instance usage |
| ml.c6in.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.12xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c6in.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.16xlarge for endpoint usage |
| ml.c6in.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.24xlarge for endpoint usage |
| ml.c6in.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.2xlarge for endpoint usage |
| ml.c6in.32xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.32xlarge for endpoint usage |
| ml.c6in.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.4xlarge for endpoint usage |
| ml.c6in.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.8xlarge for endpoint usage |
| ml.c6in.large for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.large for endpoint usage |
| ml.c6in.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c6in.xlarge for endpoint usage |
| ml.c7g.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.12xlarge for cluster spot instance usage |
| ml.c7g.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c7g.12xlarge for cluster usage |
| ml.c7g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.12xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7g.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.16xlarge for cluster spot instance usage |
| ml.c7g.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c7g.16xlarge for cluster usage |
| ml.c7g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.16xlarge for endpoint usage |
| ml.c7g.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.2xlarge for cluster spot instance usage |
| ml.c7g.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c7g.2xlarge for cluster usage |
| ml.c7g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.2xlarge for endpoint usage |
| ml.c7g.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.4xlarge for cluster spot instance usage |
| ml.c7g.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c7g.4xlarge for cluster usage |
| ml.c7g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.4xlarge for endpoint usage |
| ml.c7g.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.8xlarge for cluster spot instance usage |
| ml.c7g.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c7g.8xlarge for cluster usage |
| ml.c7g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.8xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7g.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.large for cluster spot instance usage |
| ml.c7g.large for cluster usage | Each supported Region: 0 | Yes | ml.c7g.large for cluster usage |
| ml.c7g.large for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.large for endpoint usage |
| ml.c7g.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.medium for cluster spot instance usage |
| ml.c7g.medium for cluster usage | Each supported Region: 0 | Yes | ml.c7g.medium for cluster usage |
| ml.c7g.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c7g.xlarge for cluster spot instance usage |
| ml.c7g.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c7g.xlarge for cluster usage |
| ml.c7g.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7g.xlarge for endpoint usage |
| ml.c7i.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for endpoint usage |
| ml.c7i.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for notebook instance usage |
| ml.c7i.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for processing job usage |
| ml.c7i.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7i.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for training job usage |
| ml.c7i.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for training warm pool usage |
| ml.c7i.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.12xlarge for transform job usage |
| ml.c7i.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for endpoint usage |
| ml.c7i.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for notebook instance usage |
| ml.c7i.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for processing job usage |
| ml.c7i.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for spot training job usage |
| ml.c7i.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for training job usage |
| ml.c7i.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for training warm pool usage |
| ml.c7i.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.16xlarge for transform job usage |
| ml.c7i.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7i.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for notebook instance usage |
| ml.c7i.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for processing job usage |
| ml.c7i.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for spot training job usage |
| ml.c7i.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for training job usage |
| ml.c7i.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for training warm pool usage |
| ml.c7i.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.24xlarge for transform job usage |
| ml.c7i.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for endpoint usage |
| ml.c7i.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for notebook instance usage |
| ml.c7i.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for processing job usage |
| ml.c7i.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for spot training job usage |
| ml.c7i.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for training job usage |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ml.c7i.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for training warm pool usage |
| ml.c7i.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.2xlarge for transform job usage |
| ml.c7i.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for endpoint usage |
| ml.c7i.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for notebook instance usage |
| ml.c7i.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for processing job usage |
| ml.c7i.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for spot training job usage |
| ml.c7i.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for training job usage |
| ml.c7i.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for training warm pool usage |
| ml.c7i.48xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.48xlarge for transform job usage |
| ml.c7i.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for endpoint usage |
| ml.c7i.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7i.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for processing job usage |
| ml.c7i.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for spot training job usage |
| ml.c7i.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for training job usage |
| ml.c7i.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for training warm pool usage |
| ml.c7i.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.4xlarge for transform job usage |
| ml.c7i.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for endpoint usage |
| ml.c7i.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for notebook instance usage |
| ml.c7i.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for processing job usage |
| ml.c7i.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for spot training job usage |
| ml.c7i.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for training job usage |
| ml.c7i.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7i.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.8xlarge for transform job usage |
| ml.c7i.large for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.large for endpoint usage |
| ml.c7i.large for notebook instance usage | Each supported Region: 1 | Yes | ml.c7i.large for notebook instance usage |
| ml.c7i.large for processing job usage | Each supported Region: 0 | Yes | ml.c7i.large for processing job usage |
| ml.c7i.large for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.large for spot training job usage |
| ml.c7i.large for training job usage | Each supported Region: 0 | Yes | ml.c7i.large for training job usage |
| ml.c7i.large for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.large for training warm pool usage |
| ml.c7i.large for transform job usage | Each supported Region: 0 | Yes | ml.c7i.large for transform job usage |
| ml.c7i.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c7i.xlarge for endpoint usage |
| ml.c7i.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.c7i.xlarge for notebook instance usage |
| ml.c7i.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.c7i.xlarge for processing job usage |
| ml.c7i.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.c7i.xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c7i.xlarge for training job usage | Each supported Region: 0 | Yes | ml.c7i.xlarge for training job usage |
| ml.c7i.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.c7i.xlarge for training warm pool usage |
| ml.c7i.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.c7i.xlarge for transform job usage |
| ml.c8g.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.12xlarge for cluster spot instance usage |
| ml.c8g.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.12xlarge for cluster usage |
| ml.c8g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.12xlarge for endpoint usage |
| ml.c8g.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.16xlarge for cluster spot instance usage |
| ml.c8g.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.16xlarge for cluster usage |
| ml.c8g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.16xlarge for endpoint usage |
| ml.c8g.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.24xlarge for cluster spot instance usage |
| ml.c8g.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.24xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c8g.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.24xlarge for endpoint usage |
| ml.c8g.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.2xlarge for cluster spot instance usage |
| ml.c8g.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.2xlarge for cluster usage |
| ml.c8g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.2xlarge for endpoint usage |
| ml.c8g.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.48xlarge for cluster spot instance usage |
| ml.c8g.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.48xlarge for cluster usage |
| ml.c8g.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.48xlarge for endpoint usage |
| ml.c8g.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.4xlarge for cluster spot instance usage |
| ml.c8g.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.4xlarge for cluster usage |
| ml.c8g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.4xlarge for endpoint usage |
| ml.c8g.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.8xlarge for cluster spot instance usage |
| ml.c8g.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.8xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.c8g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.8xlarge for endpoint usage |
| ml.c8g.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.large for cluster spot instance usage |
| ml.c8g.large for cluster usage | Each supported Region: 0 | Yes | ml.c8g.large for cluster usage |
| ml.c8g.large for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.large for endpoint usage |
| ml.c8g.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.medium for cluster spot instance usage |
| ml.c8g.medium for cluster usage | Each supported Region: 0 | Yes | ml.c8g.medium for cluster usage |
| ml.c8g.medium for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.medium for endpoint usage |
| ml.c8g.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.c8g.xlarge for cluster spot instance usage |
| ml.c8g.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.c8g.xlarge for cluster usage |
| ml.c8g.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.c8g.xlarge for endpoint usage |
| ml.eia1.large for endpoint usage | Each supported Region: 0 | Yes | ml.eia1.large for endpoint usage |
| ml.eia1.medium for endpoint usage | Each supported Region: 0 | Yes | ml.eia1.medium for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.eia1.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.eia1.xlarge for endpoint usage |
| ml.eia2.large for endpoint usage | Each supported Region: 0 | Yes | ml.eia2.large for endpoint usage |
| ml.eia2.medium for endpoint usage | Each supported Region: 0 | Yes | ml.eia2.medium for endpoint usage |
| ml.eia2.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.eia2.xlarge for endpoint usage |
| ml.g4dn.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for cluster usage |
| ml.g4dn.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for endpoint usage |
| ml.g4dn.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for notebook instance usage |
| ml.g4dn.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for processing job usage |
| ml.g4dn.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for spot training job usage |
| ml.g4dn.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for training job usage |
| ml.g4dn.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for training warm pool usage |
| ml.g4dn.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g4dn.12xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g4dn.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for cluster usage |
| ml.g4dn.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for endpoint usage |
| ml.g4dn.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for notebook instance usage |
| ml.g4dn.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for processing job usage |
| ml.g4dn.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for spot training job usage |
| ml.g4dn.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for training job usage |
| ml.g4dn.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for training warm pool usage |
| ml.g4dn.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g4dn.16xlarge for transform job usage |
| ml.g4dn.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for cluster usage |
| ml.g4dn.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for endpoint usage |
| ml.g4dn.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for notebook instance usage |
| ml.g4dn.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g4dn.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for spot training job usage |
| ml.g4dn.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for training job usage |
| ml.g4dn.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for training warm pool usage |
| ml.g4dn.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g4dn.2xlarge for transform job usage |
| ml.g4dn.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for cluster usage |
| ml.g4dn.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for endpoint usage |
| ml.g4dn.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for notebook instance usage |
| ml.g4dn.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for processing job usage |
| ml.g4dn.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for spot training job usage |
| ml.g4dn.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for training job usage |
| ml.g4dn.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g4dn.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g4dn.4xlarge for transform job usage |
| ml.g4dn.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for cluster usage |
| ml.g4dn.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for endpoint usage |
| ml.g4dn.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for notebook instance usage |
| ml.g4dn.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for processing job usage |
| ml.g4dn.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for spot training job usage |
| ml.g4dn.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for training job usage |
| ml.g4dn.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for training warm pool usage |
| ml.g4dn.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g4dn.8xlarge for transform job usage |
| ml.g4dn.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for cluster usage |
| ml.g4dn.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for endpoint usage |
| ml.g4dn.xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g4dn.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for processing job usage |
| ml.g4dn.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for spot training job usage |
| ml.g4dn.xlarge for training job usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for training job usage |
| ml.g4dn.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for training warm pool usage |
| ml.g4dn.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g4dn.xlarge for transform job usage |
| ml.g5.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for cluster spot instance usage |
| ml.g5.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for cluster usage |
| ml.g5.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for endpoint usage |
| ml.g5.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for notebook instance usage |
| ml.g5.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for processing job usage |
| ml.g5.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g5.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for training job usage |
| ml.g5.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for training warm pool usage |
| ml.g5.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.12xlarge for transform job usage |
| ml.g5.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for cluster spot instance usage |
| ml.g5.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for cluster usage |
| ml.g5.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for endpoint usage |
| ml.g5.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for notebook instance usage |
| ml.g5.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for processing job usage |
| ml.g5.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for spot training job usage |
| ml.g5.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for training job usage |
| ml.g5.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g5.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.16xlarge for transform job usage |
| ml.g5.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for cluster spot instance usage |
| ml.g5.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for cluster usage |
| ml.g5.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for endpoint usage |
| ml.g5.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for notebook instance usage |
| ml.g5.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for processing job usage |
| ml.g5.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for spot training job usage |
| ml.g5.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for training job usage |
| ml.g5.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for training warm pool usage |
| ml.g5.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.24xlarge for transform job usage |
| ml.g5.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g5.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for cluster usage |
| ml.g5.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for endpoint usage |
| ml.g5.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for notebook instance usage |
| ml.g5.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for processing job usage |
| ml.g5.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for spot training job usage |
| ml.g5.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for training job usage |
| ml.g5.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for training warm pool usage |
| ml.g5.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.2xlarge for transform job usage |
| ml.g5.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for cluster spot instance usage |
| ml.g5.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for cluster usage |
| ml.g5.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g5.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for notebook instance usage |
| ml.g5.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for processing job usage |
| ml.g5.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for spot training job usage |
| ml.g5.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for training job usage |
| ml.g5.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for training warm pool usage |
| ml.g5.48xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.48xlarge for transform job usage |
| ml.g5.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for cluster spot instance usage |
| ml.g5.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for cluster usage |
| ml.g5.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for endpoint usage |
| ml.g5.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for notebook instance usage |
| ml.g5.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for processing job usage |
| ml.g5.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g5.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for training job usage |
| ml.g5.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for training warm pool usage |
| ml.g5.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.4xlarge for transform job usage |
| ml.g5.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for cluster spot instance usage |
| ml.g5.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for cluster usage |
| ml.g5.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for endpoint usage |
| ml.g5.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for notebook instance usage |
| ml.g5.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for processing job usage |
| ml.g5.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for spot training job usage |
| ml.g5.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for training job usage |
| ml.g5.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g5.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.8xlarge for transform job usage |
| ml.g5.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g5.xlarge for cluster spot instance usage |
| ml.g5.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g5.xlarge for cluster usage |
| ml.g5.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g5.xlarge for endpoint usage |
| ml.g5.xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g5.xlarge for notebook instance usage |
| ml.g5.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g5.xlarge for processing job usage |
| ml.g5.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g5.xlarge for spot training job usage |
| ml.g5.xlarge for training job usage | Each supported Region: 0 | Yes | ml.g5.xlarge for training job usage |
| ml.g5.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g5.xlarge for training warm pool usage |
| ml.g5.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.g5.xlarge for transform job usage |
| ml.g6.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for cluster spot instance usage |
| ml.g6.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for endpoint usage |
| ml.g6.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for notebook instance usage |
| ml.g6.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for processing job usage |
| ml.g6.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for spot training job usage |
| ml.g6.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for training job usage |
| ml.g6.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.12xlarge for training warm pool usage |
| ml.g6.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for cluster spot instance usage |
| ml.g6.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for cluster usage |
| ml.g6.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for endpoint usage |
| ml.g6.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for notebook instance usage |
| ml.g6.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for spot training job usage |
| ml.g6.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for training job usage |
| ml.g6.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.16xlarge for training warm pool usage |
| ml.g6.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for cluster spot instance usage |
| ml.g6.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for cluster usage |
| ml.g6.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for endpoint usage |
| ml.g6.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for notebook instance usage |
| ml.g6.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for processing job usage |
| ml.g6.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for spot training job usage |
| ml.g6.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for training job usage |
| ml.g6.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.24xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for cluster spot instance usage |
| ml.g6.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for cluster usage |
| ml.g6.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for endpoint usage |
| ml.g6.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for notebook instance usage |
| ml.g6.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for processing job usage |
| ml.g6.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for spot training job usage |
| ml.g6.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for training job usage |
| ml.g6.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.2xlarge for training warm pool usage |
| ml.g6.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for cluster spot instance usage |
| ml.g6.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for cluster usage |
| ml.g6.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for notebook instance usage |
| ml.g6.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for processing job usage |
| ml.g6.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for spot training job usage |
| ml.g6.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for training job usage |
| ml.g6.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.48xlarge for training warm pool usage |
| ml.g6.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for cluster spot instance usage |
| ml.g6.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for cluster usage |
| ml.g6.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for endpoint usage |
| ml.g6.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for notebook instance usage |
| ml.g6.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for processing job usage |
| ml.g6.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for spot training job usage |
| ml.g6.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.4xlarge for training warm pool usage |
| ml.g6.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for cluster spot instance usage |
| ml.g6.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for cluster usage |
| ml.g6.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for endpoint usage |
| ml.g6.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for notebook instance usage |
| ml.g6.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for processing job usage |
| ml.g6.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for spot training job usage |
| ml.g6.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for training job usage |
| ml.g6.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.8xlarge for training warm pool usage |
| ml.g6.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6.xlarge for cluster spot instance usage |
| ml.g6.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6.xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6.xlarge for endpoint usage |
| ml.g6.xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.g6.xlarge for notebook instance usage |
| ml.g6.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6.xlarge for processing job usage |
| ml.g6.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6.xlarge for spot training job usage |
| ml.g6.xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6.xlarge for training job usage |
| ml.g6.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6.xlarge for training warm pool usage |
| ml.g6e.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for cluster spot instance usage |
| ml.g6e.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for cluster usage |
| ml.g6e.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for endpoint usage |
| ml.g6e.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for processing job usage |
| ml.g6e.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for spot training job usage |
| ml.g6e.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6e.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.12xlarge for training warm pool usage |
| ml.g6e.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for cluster spot instance usage |
| ml.g6e.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for cluster usage |
| ml.g6e.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for endpoint usage |
| ml.g6e.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for processing job usage |
| ml.g6e.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for spot training job usage |
| ml.g6e.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for training job usage |
| ml.g6e.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.16xlarge for training warm pool usage |
| ml.g6e.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for cluster spot instance usage |
| ml.g6e.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for cluster usage |
| ml.g6e.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6e.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for processing job usage |
| ml.g6e.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for spot training job usage |
| ml.g6e.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for training job usage |
| ml.g6e.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.24xlarge for training warm pool usage |
| ml.g6e.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for cluster spot instance usage |
| ml.g6e.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for cluster usage |
| ml.g6e.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for endpoint usage |
| ml.g6e.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for processing job usage |
| ml.g6e.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for spot training job usage |
| ml.g6e.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for training job usage |
| ml.g6e.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.2xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6e.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for cluster spot instance usage |
| ml.g6e.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for cluster usage |
| ml.g6e.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for endpoint usage |
| ml.g6e.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for processing job usage |
| ml.g6e.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for spot training job usage |
| ml.g6e.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for training job usage |
| ml.g6e.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.48xlarge for training warm pool usage |
| ml.g6e.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for cluster spot instance usage |
| ml.g6e.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for cluster usage |
| ml.g6e.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for endpoint usage |
| ml.g6e.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6e.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for spot training job usage |
| ml.g6e.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for training job usage |
| ml.g6e.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.4xlarge for training warm pool usage |
| ml.g6e.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for cluster spot instance usage |
| ml.g6e.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for cluster usage |
| ml.g6e.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for endpoint usage |
| ml.g6e.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for processing job usage |
| ml.g6e.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for spot training job usage |
| ml.g6e.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for training job usage |
| ml.g6e.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.8xlarge for training warm pool usage |
| ml.g6e.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g6e.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for cluster usage |
| ml.g6e.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for endpoint usage |
| ml.g6e.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for processing job usage |
| ml.g6e.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for spot training job usage |
| ml.g6e.xlarge for training job usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for training job usage |
| ml.g6e.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g6e.xlarge for training warm pool usage |
| ml.g7.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for cluster spot instance usage |
| ml.g7.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for cluster usage |
| ml.g7.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for endpoint usage |
| ml.g7.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for processing job usage |
| ml.g7.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for spot training job usage |
| ml.g7.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7.12xlarge for training warm pool usage |
| ml.g7.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.16xlarge for endpoint usage |
| ml.g7.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for cluster spot instance usage |
| ml.g7.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for cluster usage |
| ml.g7.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for endpoint usage |
| ml.g7.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for processing job usage |
| ml.g7.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for spot training job usage |
| ml.g7.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for training job usage |
| ml.g7.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7.24xlarge for training warm pool usage |
| ml.g7.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for cluster spot instance usage |
| ml.g7.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for endpoint usage |
| ml.g7.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for processing job usage |
| ml.g7.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for spot training job usage |
| ml.g7.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for training job usage |
| ml.g7.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7.2xlarge for training warm pool usage |
| ml.g7.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for cluster spot instance usage |
| ml.g7.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for cluster usage |
| ml.g7.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for endpoint usage |
| ml.g7.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for processing job usage |
| ml.g7.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for spot training job usage |
| ml.g7.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7.48xlarge for training warm pool usage |
| ml.g7.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for cluster spot instance usage |
| ml.g7.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for cluster usage |
| ml.g7.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for endpoint usage |
| ml.g7.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for processing job usage |
| ml.g7.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for spot training job usage |
| ml.g7.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for training job usage |
| ml.g7.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7.4xlarge for training warm pool usage |
| ml.g7.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for cluster spot instance usage |
| ml.g7.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for cluster usage |
| ml.g7.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for processing job usage |
| ml.g7.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for spot training job usage |
| ml.g7.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for training job usage |
| ml.g7.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7.8xlarge for training warm pool usage |
| ml.g7e.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for cluster spot instance usage |
| ml.g7e.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for cluster usage |
| ml.g7e.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for endpoint usage |
| ml.g7e.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for processing job usage |
| ml.g7e.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for spot training job usage |
| ml.g7e.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for training job usage |
| ml.g7e.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7e.12xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7e.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for cluster spot instance usage |
| ml.g7e.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for cluster usage |
| ml.g7e.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for endpoint usage |
| ml.g7e.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for processing job usage |
| ml.g7e.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for spot training job usage |
| ml.g7e.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for training job usage |
| ml.g7e.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7e.24xlarge for training warm pool usage |
| ml.g7e.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for cluster spot instance usage |
| ml.g7e.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for cluster usage |
| ml.g7e.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for endpoint usage |
| ml.g7e.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7e.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for spot training job usage |
| ml.g7e.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for training job usage |
| ml.g7e.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7e.2xlarge for training warm pool usage |
| ml.g7e.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for cluster spot instance usage |
| ml.g7e.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for cluster usage |
| ml.g7e.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for endpoint usage |
| ml.g7e.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for processing job usage |
| ml.g7e.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for spot training job usage |
| ml.g7e.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for training job usage |
| ml.g7e.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7e.48xlarge for training warm pool usage |
| ml.g7e.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7e.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for cluster usage |
| ml.g7e.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for endpoint usage |
| ml.g7e.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for processing job usage |
| ml.g7e.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for spot training job usage |
| ml.g7e.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for training job usage |
| ml.g7e.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7e.4xlarge for training warm pool usage |
| ml.g7e.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for cluster spot instance usage |
| ml.g7e.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for cluster usage |
| ml.g7e.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for endpoint usage |
| ml.g7e.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for processing job usage |
| ml.g7e.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.g7e.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for training job usage |
| ml.g7e.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.g7e.8xlarge for training warm pool usage |
| ml.gr6.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.gr6.4xlarge for cluster spot instance usage |
| ml.gr6.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.gr6.4xlarge for cluster usage |
| ml.gr6.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.gr6.8xlarge for cluster spot instance usage |
| ml.gr6.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.gr6.8xlarge for cluster usage |
| ml.i3en.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.12xlarge for cluster spot instance usage |
| ml.i3en.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.12xlarge for cluster usage |
| ml.i3en.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.16xlarge for cluster spot instance usage |
| ml.i3en.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.16xlarge for cluster usage |
| ml.i3en.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.24xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.i3en.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.24xlarge for cluster usage |
| ml.i3en.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.2xlarge for cluster spot instance usage |
| ml.i3en.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.2xlarge for cluster usage |
| ml.i3en.3xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.3xlarge for cluster spot instance usage |
| ml.i3en.3xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.3xlarge for cluster usage |
| ml.i3en.6xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.6xlarge for cluster spot instance usage |
| ml.i3en.6xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.6xlarge for cluster usage |
| ml.i3en.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.large for cluster spot instance usage |
| ml.i3en.large for cluster usage | Each supported Region: 0 | Yes | ml.i3en.large for cluster usage |
| ml.i3en.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.i3en.xlarge for cluster spot instance usage |
| ml.i3en.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.i3en.xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.inf1.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf1.24xlarge for endpoint usage |
| ml.inf1.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf1.24xlarge for notebook instance usage |
| ml.inf1.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf1.2xlarge for endpoint usage |
| ml.inf1.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf1.2xlarge for notebook instance usage |
| ml.inf1.6xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf1.6xlarge for endpoint usage |
| ml.inf1.6xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf1.6xlarge for notebook instance usage |
| ml.inf1.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf1.xlarge for endpoint usage |
| ml.inf1.xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf1.xlarge for notebook instance usage |
| ml.inf2.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf2.24xlarge for endpoint usage |
| ml.inf2.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf2.24xlarge for notebook instance usage |
| ml.inf2.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.inf2.24xlarge for transform job usage |
| ml.inf2.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf2.48xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.inf2.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf2.48xlarge for notebook instance usage |
| ml.inf2.48xlarge for transform job usage | Each supported Region: 0 | Yes | ml.inf2.48xlarge for transform job usage |
| ml.inf2.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf2.8xlarge for endpoint usage |
| ml.inf2.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf2.8xlarge for notebook instance usage |
| ml.inf2.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.inf2.8xlarge for transform job usage |
| ml.inf2.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf2.xlarge for endpoint usage |
| ml.inf2.xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.inf2.xlarge for notebook instance usage |
| ml.inf2.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.inf2.xlarge for transform job usage |
| ml.inf2e.32xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.inf2e.32xlarge for endpoint usage |
| ml.m4.10xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for endpoint usage |
| ml.m4.10xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for notebook instance usage |
| ml.m4.10xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m4.10xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for spot training job usage |
| ml.m4.10xlarge for training job usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for training job usage |
| ml.m4.10xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for training warm pool usage |
| ml.m4.10xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m4.10xlarge for transform job usage |
| ml.m4.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for endpoint usage |
| ml.m4.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for notebook instance usage |
| ml.m4.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for processing job usage |
| ml.m4.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for spot training job usage |
| ml.m4.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for training job usage |
| ml.m4.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for training warm pool usage |
| ml.m4.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m4.16xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m4.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m4.2xlarge for endpoint usage |
| ml.m4.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m4.2xlarge for notebook instance usage |
| ml.m4.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m4.2xlarge for processing job usage |
| ml.m4.2xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.m4.2xlarge for spot training job usage |
| ml.m4.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.m4.2xlarge for training job usage |
| ml.m4.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m4.2xlarge for training warm pool usage |
| ml.m4.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m4.2xlarge for transform job usage |
| ml.m4.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m4.4xlarge for endpoint usage |
| ml.m4.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m4.4xlarge for notebook instance usage |
| ml.m4.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m4.4xlarge for processing job usage |
| ml.m4.4xlarge for spot training job usage | Each supported Region: 2 | Yes | ml.m4.4xlarge for spot training job usage |
| ml.m4.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.m4.4xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m4.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m4.4xlarge for training warm pool usage |
| ml.m4.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m4.4xlarge for transform job usage |
| ml.m4.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m4.xlarge for endpoint usage |
| ml.m4.xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m4.xlarge for notebook instance usage |
| ml.m4.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m4.xlarge for processing job usage |
| ml.m4.xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.m4.xlarge for spot training job usage |
| ml.m4.xlarge for training job usage | Each supported Region: 0 | Yes | ml.m4.xlarge for training job usage |
| ml.m4.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m4.xlarge for training warm pool usage |
| ml.m4.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m4.xlarge for transform job usage |
| ml.m5.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for cluster spot instance usage |
| ml.m5.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m5.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for endpoint usage |
| ml.m5.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for notebook instance usage |
| ml.m5.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for processing job usage |
| ml.m5.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for spot training job usage |
| ml.m5.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for training job usage |
| ml.m5.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for training warm pool usage |
| ml.m5.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m5.12xlarge for transform job usage |
| ml.m5.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.16xlarge for cluster spot instance usage |
| ml.m5.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.16xlarge for cluster usage |
| ml.m5.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for cluster spot instance usage |
| ml.m5.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m5.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for endpoint usage |
| ml.m5.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for notebook instance usage |
| ml.m5.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for processing job usage |
| ml.m5.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for spot training job usage |
| ml.m5.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for training job usage |
| ml.m5.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for training warm pool usage |
| ml.m5.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m5.24xlarge for transform job usage |
| ml.m5.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for cluster spot instance usage |
| ml.m5.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for cluster usage |
| ml.m5.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for endpoint usage |
| ml.m5.2xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.m5.2xlarge for notebook instance usage |
| ml.m5.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m5.2xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.m5.2xlarge for spot training job usage |
| ml.m5.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for training job usage |
| ml.m5.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for training warm pool usage |
| ml.m5.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m5.2xlarge for transform job usage |
| ml.m5.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for cluster spot instance usage |
| ml.m5.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for cluster usage |
| ml.m5.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for endpoint usage |
| ml.m5.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for notebook instance usage |
| ml.m5.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for processing job usage |
| ml.m5.4xlarge for spot training job usage | Each supported Region: 2 | Yes | ml.m5.4xlarge for spot training job usage |
| ml.m5.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m5.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for training warm pool usage |
| ml.m5.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m5.4xlarge for transform job usage |
| ml.m5.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.8xlarge for cluster spot instance usage |
| ml.m5.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.8xlarge for cluster usage |
| ml.m5.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.large for cluster spot instance usage |
| ml.m5.large for cluster usage | Each supported Region: 0 | Yes | ml.m5.large for cluster usage |
| ml.m5.large for endpoint usage | Each supported Region: 0 | Yes | ml.m5.large for endpoint usage |
| ml.m5.large for processing job usage | Each supported Region: 0 | Yes | ml.m5.large for processing job usage |
| ml.m5.large for spot training job usage | Each supported Region: 4 | Yes | ml.m5.large for spot training job usage |
| ml.m5.large for training job usage | Each supported Region: 0 | Yes | ml.m5.large for training job usage |
| ml.m5.large for training warm pool usage | Each supported Region: 0 | Yes | ml.m5.large for training warm pool usage |
| ml.m5.large for transform job usage | Each supported Region: 0 | Yes | ml.m5.large for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m5.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m5.xlarge for cluster spot instance usage |
| ml.m5.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m5.xlarge for cluster usage |
| ml.m5.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5.xlarge for endpoint usage |
| ml.m5.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.m5.xlarge for notebook instance usage |
| ml.m5.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m5.xlarge for processing job usage |
| ml.m5.xlarge for spot training job usage | Each supported Region: 4 | Yes | ml.m5.xlarge for spot training job usage |
| ml.m5.xlarge for training job usage | Each supported Region: 0 | Yes | ml.m5.xlarge for training job usage |
| ml.m5.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m5.xlarge for training warm pool usage |
| ml.m5.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m5.xlarge for transform job usage |
| ml.m5d.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5d.12xlarge for endpoint usage |
| ml.m5d.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5d.12xlarge for notebook instance usage |
| ml.m5d.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5d.16xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m5d.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5d.24xlarge for endpoint usage |
| ml.m5d.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5d.24xlarge for notebook instance usage |
| ml.m5d.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5d.2xlarge for endpoint usage |
| ml.m5d.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5d.2xlarge for notebook instance usage |
| ml.m5d.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5d.4xlarge for endpoint usage |
| ml.m5d.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5d.4xlarge for notebook instance usage |
| ml.m5d.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m5d.8xlarge for notebook instance usage |
| ml.m5d.large for endpoint usage | Each supported Region: 0 | Yes | ml.m5d.large for endpoint usage |
| ml.m5d.large for notebook instance usage | Each supported Region: 1 | Yes | ml.m5d.large for notebook instance usage |
| ml.m5d.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m5d.xlarge for endpoint usage |
| ml.m5d.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.m5d.xlarge for notebook instance usage |
| ml.m6a.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.12xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6a.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.12xlarge for cluster usage |
| ml.m6a.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.16xlarge for cluster spot instance usage |
| ml.m6a.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.16xlarge for cluster usage |
| ml.m6a.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.24xlarge for cluster spot instance usage |
| ml.m6a.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.24xlarge for cluster usage |
| ml.m6a.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.2xlarge for cluster spot instance usage |
| ml.m6a.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.2xlarge for cluster usage |
| ml.m6a.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.32xlarge for cluster spot instance usage |
| ml.m6a.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.32xlarge for cluster usage |
| ml.m6a.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.48xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6a.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.48xlarge for cluster usage |
| ml.m6a.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.4xlarge for cluster spot instance usage |
| ml.m6a.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.4xlarge for cluster usage |
| ml.m6a.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.8xlarge for cluster spot instance usage |
| ml.m6a.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.8xlarge for cluster usage |
| ml.m6a.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.large for cluster spot instance usage |
| ml.m6a.large for cluster usage | Each supported Region: 0 | Yes | ml.m6a.large for cluster usage |
| ml.m6a.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6a.xlarge for cluster spot instance usage |
| ml.m6a.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6a.xlarge for cluster usage |
| ml.m6g.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.12xlarge for cluster spot instance usage |
| ml.m6g.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6g.12xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6g.12xlarge for endpoint usage |
| ml.m6g.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.16xlarge for cluster spot instance usage |
| ml.m6g.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6g.16xlarge for cluster usage |
| ml.m6g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6g.16xlarge for endpoint usage |
| ml.m6g.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.2xlarge for cluster spot instance usage |
| ml.m6g.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6g.2xlarge for cluster usage |
| ml.m6g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6g.2xlarge for endpoint usage |
| ml.m6g.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.4xlarge for cluster spot instance usage |
| ml.m6g.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6g.4xlarge for cluster usage |
| ml.m6g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6g.4xlarge for endpoint usage |
| ml.m6g.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.8xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6g.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6g.8xlarge for cluster usage |
| ml.m6g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6g.8xlarge for endpoint usage |
| ml.m6g.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.large for cluster spot instance usage |
| ml.m6g.large for cluster usage | Each supported Region: 0 | Yes | ml.m6g.large for cluster usage |
| ml.m6g.large for endpoint usage | Each supported Region: 2 | Yes | ml.m6g.large for endpoint usage |
| ml.m6g.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.medium for cluster spot instance usage |
| ml.m6g.medium for cluster usage | Each supported Region: 0 | Yes | ml.m6g.medium for cluster usage |
| ml.m6g.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6g.xlarge for cluster spot instance usage |
| ml.m6g.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6g.xlarge for cluster usage |
| ml.m6g.xlarge for endpoint usage | Each supported Region: 1 | Yes | ml.m6g.xlarge for endpoint usage |
| ml.m6gd.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.12xlarge for endpoint usage |
| ml.m6gd.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.16xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6gd.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.2xlarge for endpoint usage |
| ml.m6gd.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.4xlarge for endpoint usage |
| ml.m6gd.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.8xlarge for endpoint usage |
| ml.m6gd.large for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.large for endpoint usage |
| ml.m6gd.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6gd.xlarge for endpoint usage |
| ml.m6i.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for cluster spot instance usage |
| ml.m6i.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for cluster usage |
| ml.m6i.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for endpoint usage |
| ml.m6i.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for notebook instance usage |
| ml.m6i.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for processing job usage |
| ml.m6i.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for spot training job usage |
| ml.m6i.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for training warm pool usage |
| ml.m6i.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.12xlarge for transform job usage |
| ml.m6i.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for cluster spot instance usage |
| ml.m6i.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for cluster usage |
| ml.m6i.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for endpoint usage |
| ml.m6i.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for notebook instance usage |
| ml.m6i.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for processing job usage |
| ml.m6i.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for spot training job usage |
| ml.m6i.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for training job usage |
| ml.m6i.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for training warm pool usage |
| ml.m6i.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.16xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for cluster spot instance usage |
| ml.m6i.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for cluster usage |
| ml.m6i.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for endpoint usage |
| ml.m6i.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for notebook instance usage |
| ml.m6i.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for processing job usage |
| ml.m6i.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for spot training job usage |
| ml.m6i.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for training job usage |
| ml.m6i.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for training warm pool usage |
| ml.m6i.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.24xlarge for transform job usage |
| ml.m6i.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for cluster spot instance usage |
| ml.m6i.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for endpoint usage |
| ml.m6i.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for notebook instance usage |
| ml.m6i.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for processing job usage |
| ml.m6i.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for spot training job usage |
| ml.m6i.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for training job usage |
| ml.m6i.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for training warm pool usage |
| ml.m6i.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.2xlarge for transform job usage |
| ml.m6i.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for cluster spot instance usage |
| ml.m6i.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for cluster usage |
| ml.m6i.32xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for endpoint usage |
| ml.m6i.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.32xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for processing job usage |
| ml.m6i.32xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for spot training job usage |
| ml.m6i.32xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for training job usage |
| ml.m6i.32xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for training warm pool usage |
| ml.m6i.32xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.32xlarge for transform job usage |
| ml.m6i.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for cluster spot instance usage |
| ml.m6i.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for cluster usage |
| ml.m6i.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for endpoint usage |
| ml.m6i.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for notebook instance usage |
| ml.m6i.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for processing job usage |
| ml.m6i.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for training job usage |
| ml.m6i.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for training warm pool usage |
| ml.m6i.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.4xlarge for transform job usage |
| ml.m6i.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for cluster spot instance usage |
| ml.m6i.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for cluster usage |
| ml.m6i.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for endpoint usage |
| ml.m6i.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for notebook instance usage |
| ml.m6i.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for processing job usage |
| ml.m6i.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for spot training job usage |
| ml.m6i.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for training job usage |
| ml.m6i.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.8xlarge for transform job usage |
| ml.m6i.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.large for cluster spot instance usage |
| ml.m6i.large for cluster usage | Each supported Region: 0 | Yes | ml.m6i.large for cluster usage |
| ml.m6i.large for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.large for endpoint usage |
| ml.m6i.large for notebook instance usage | Each supported Region: 1 | Yes | ml.m6i.large for notebook instance usage |
| ml.m6i.large for processing job usage | Each supported Region: 0 | Yes | ml.m6i.large for processing job usage |
| ml.m6i.large for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.large for spot training job usage |
| ml.m6i.large for training job usage | Each supported Region: 0 | Yes | ml.m6i.large for training job usage |
| ml.m6i.large for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.large for training warm pool usage |
| ml.m6i.large for transform job usage | Each supported Region: 0 | Yes | ml.m6i.large for transform job usage |
| ml.m6i.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for cluster spot instance usage |
| ml.m6i.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6i.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for endpoint usage |
| ml.m6i.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.m6i.xlarge for notebook instance usage |
| ml.m6i.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for processing job usage |
| ml.m6i.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for spot training job usage |
| ml.m6i.xlarge for training job usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for training job usage |
| ml.m6i.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for training warm pool usage |
| ml.m6i.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m6i.xlarge for transform job usage |
| ml.m6id.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.12xlarge for notebook instance usage |
| ml.m6id.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.16xlarge for notebook instance usage |
| ml.m6id.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.24xlarge for notebook instance usage |
| ml.m6id.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.2xlarge for notebook instance usage |
| ml.m6id.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.32xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m6id.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.4xlarge for notebook instance usage |
| ml.m6id.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m6id.8xlarge for notebook instance usage |
| ml.m6id.large for notebook instance usage | Each supported Region: 1 | Yes | ml.m6id.large for notebook instance usage |
| ml.m6id.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.m6id.xlarge for notebook instance usage |
| ml.m7g.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.12xlarge for cluster spot instance usage |
| ml.m7g.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7g.12xlarge for cluster usage |
| ml.m7g.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.16xlarge for cluster spot instance usage |
| ml.m7g.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7g.16xlarge for cluster usage |
| ml.m7g.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.2xlarge for cluster spot instance usage |
| ml.m7g.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7g.2xlarge for cluster usage |
| ml.m7g.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.4xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7g.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7g.4xlarge for cluster usage |
| ml.m7g.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.8xlarge for cluster spot instance usage |
| ml.m7g.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7g.8xlarge for cluster usage |
| ml.m7g.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.large for cluster spot instance usage |
| ml.m7g.large for cluster usage | Each supported Region: 0 | Yes | ml.m7g.large for cluster usage |
| ml.m7g.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.medium for cluster spot instance usage |
| ml.m7g.medium for cluster usage | Each supported Region: 0 | Yes | ml.m7g.medium for cluster usage |
| ml.m7g.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7g.xlarge for cluster spot instance usage |
| ml.m7g.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7g.xlarge for cluster usage |
| ml.m7i.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for cluster spot instance usage |
| ml.m7i.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for endpoint usage |
| ml.m7i.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for notebook instance usage |
| ml.m7i.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for processing job usage |
| ml.m7i.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for spot training job usage |
| ml.m7i.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for training job usage |
| ml.m7i.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for training warm pool usage |
| ml.m7i.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.12xlarge for transform job usage |
| ml.m7i.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for cluster spot instance usage |
| ml.m7i.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for cluster usage |
| ml.m7i.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for endpoint usage |
| ml.m7i.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for processing job usage |
| ml.m7i.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for spot training job usage |
| ml.m7i.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for training job usage |
| ml.m7i.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for training warm pool usage |
| ml.m7i.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.16xlarge for transform job usage |
| ml.m7i.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for cluster spot instance usage |
| ml.m7i.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for cluster usage |
| ml.m7i.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for endpoint usage |
| ml.m7i.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for notebook instance usage |
| ml.m7i.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for processing job usage |
| ml.m7i.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for training job usage |
| ml.m7i.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for training warm pool usage |
| ml.m7i.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.24xlarge for transform job usage |
| ml.m7i.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for cluster spot instance usage |
| ml.m7i.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for cluster usage |
| ml.m7i.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for endpoint usage |
| ml.m7i.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for notebook instance usage |
| ml.m7i.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for processing job usage |
| ml.m7i.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for spot training job usage |
| ml.m7i.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for training job usage |
| ml.m7i.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.2xlarge for transform job usage |
| ml.m7i.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for cluster spot instance usage |
| ml.m7i.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for cluster usage |
| ml.m7i.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for endpoint usage |
| ml.m7i.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for notebook instance usage |
| ml.m7i.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for processing job usage |
| ml.m7i.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for spot training job usage |
| ml.m7i.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for training job usage |
| ml.m7i.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for training warm pool usage |
| ml.m7i.48xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.48xlarge for transform job usage |
| ml.m7i.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for cluster usage |
| ml.m7i.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for endpoint usage |
| ml.m7i.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for notebook instance usage |
| ml.m7i.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for processing job usage |
| ml.m7i.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for spot training job usage |
| ml.m7i.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for training job usage |
| ml.m7i.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for training warm pool usage |
| ml.m7i.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.4xlarge for transform job usage |
| ml.m7i.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for cluster spot instance usage |
| ml.m7i.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for cluster usage |
| ml.m7i.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for notebook instance usage |
| ml.m7i.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for processing job usage |
| ml.m7i.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for spot training job usage |
| ml.m7i.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for training job usage |
| ml.m7i.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for training warm pool usage |
| ml.m7i.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.8xlarge for transform job usage |
| ml.m7i.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.large for cluster spot instance usage |
| ml.m7i.large for cluster usage | Each supported Region: 0 | Yes | ml.m7i.large for cluster usage |
| ml.m7i.large for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.large for endpoint usage |
| ml.m7i.large for notebook instance usage | Each supported Region: 1 | Yes | ml.m7i.large for notebook instance usage |
| ml.m7i.large for processing job usage | Each supported Region: 0 | Yes | ml.m7i.large for processing job usage |
| ml.m7i.large for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.large for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m7i.large for training job usage | Each supported Region: 0 | Yes | ml.m7i.large for training job usage |
| ml.m7i.large for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.large for training warm pool usage |
| ml.m7i.large for transform job usage | Each supported Region: 0 | Yes | ml.m7i.large for transform job usage |
| ml.m7i.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for cluster spot instance usage |
| ml.m7i.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for cluster usage |
| ml.m7i.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for endpoint usage |
| ml.m7i.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.m7i.xlarge for notebook instance usage |
| ml.m7i.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for processing job usage |
| ml.m7i.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for spot training job usage |
| ml.m7i.xlarge for training job usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for training job usage |
| ml.m7i.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for training warm pool usage |
| ml.m7i.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.m7i.xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m8g.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.12xlarge for cluster spot instance usage |
| ml.m8g.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.12xlarge for cluster usage |
| ml.m8g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.12xlarge for endpoint usage |
| ml.m8g.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.16xlarge for cluster spot instance usage |
| ml.m8g.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.16xlarge for cluster usage |
| ml.m8g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.16xlarge for endpoint usage |
| ml.m8g.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.24xlarge for cluster spot instance usage |
| ml.m8g.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.24xlarge for cluster usage |
| ml.m8g.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.24xlarge for endpoint usage |
| ml.m8g.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.2xlarge for cluster spot instance usage |
| ml.m8g.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.2xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m8g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.2xlarge for endpoint usage |
| ml.m8g.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.48xlarge for cluster spot instance usage |
| ml.m8g.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.48xlarge for cluster usage |
| ml.m8g.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.48xlarge for endpoint usage |
| ml.m8g.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.4xlarge for cluster spot instance usage |
| ml.m8g.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.4xlarge for cluster usage |
| ml.m8g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.4xlarge for endpoint usage |
| ml.m8g.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.8xlarge for cluster spot instance usage |
| ml.m8g.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.8xlarge for cluster usage |
| ml.m8g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.8xlarge for endpoint usage |
| ml.m8g.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.large for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.m8g.large for cluster usage | Each supported Region: 0 | Yes | ml.m8g.large for cluster usage |
| ml.m8g.large for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.large for endpoint usage |
| ml.m8g.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.medium for cluster spot instance usage |
| ml.m8g.medium for cluster usage | Each supported Region: 0 | Yes | ml.m8g.medium for cluster usage |
| ml.m8g.medium for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.medium for endpoint usage |
| ml.m8g.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.m8g.xlarge for cluster spot instance usage |
| ml.m8g.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.m8g.xlarge for cluster usage |
| ml.m8g.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.m8g.xlarge for endpoint usage |
| ml.p3.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p3.16xlarge for endpoint usage |
| ml.p3.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p3.16xlarge for notebook instance usage |
| ml.p3.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.p3.16xlarge for transform job usage |
| ml.p3.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p3.2xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.p3.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p3.2xlarge for notebook instance usage |
| ml.p3.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.p3.2xlarge for transform job usage |
| ml.p3.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p3.8xlarge for endpoint usage |
| ml.p3.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p3.8xlarge for notebook instance usage |
| ml.p3.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.p3.8xlarge for transform job usage |
| ml.p3dn.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p3dn.24xlarge for notebook instance usage |
| ml.p3dn.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p3dn.24xlarge for spot training job usage |
| ml.p3dn.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.p3dn.24xlarge for training job usage |
| ml.p3dn.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p3dn.24xlarge for training warm pool usage |
| ml.p4d.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for cluster spot instance usage |
| ml.p4d.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.p4d.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for endpoint usage |
| ml.p4d.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for notebook instance usage |
| ml.p4d.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for spot training job usage |
| ml.p4d.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for training job usage |
| ml.p4d.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p4d.24xlarge for training warm pool usage |
| ml.p4de.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for cluster spot instance usage |
| ml.p4de.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for cluster usage |
| ml.p4de.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for endpoint usage |
| ml.p4de.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for notebook instance usage |
| ml.p4de.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for spot training job usage |
| ml.p4de.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.p4de.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p4de.24xlarge for training warm pool usage |
| ml.p5.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for cluster spot instance usage |
| ml.p5.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for cluster usage |
| ml.p5.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for endpoint usage |
| ml.p5.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for notebook instance usage |
| ml.p5.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for spot training job usage |
| ml.p5.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for training job usage |
| ml.p5.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p5.48xlarge for training warm pool usage |
| ml.p5.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for cluster spot instance usage |
| ml.p5.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for cluster usage |
| ml.p5.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ml.p5.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for notebook instance usage |
| ml.p5.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for processing job usage |
| ml.p5.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for spot training job usage |
| ml.p5.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for training job usage |
| ml.p5.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p5.4xlarge for training warm pool usage |
| ml.p5e.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p5e.48xlarge for cluster spot instance usage |
| ml.p5e.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p5e.48xlarge for cluster usage |
| ml.p5e.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p5e.48xlarge for endpoint usage |
| ml.p5e.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p5e.48xlarge for spot training job usage |
| ml.p5e.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.p5e.48xlarge for training job usage |
| ml.p5e.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p5e.48xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.p5en.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for cluster spot instance usage |
| ml.p5en.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for cluster usage |
| ml.p5en.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for endpoint usage |
| ml.p5en.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for notebook instance usage |
| ml.p5en.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for spot training job usage |
| ml.p5en.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for training job usage |
| ml.p5en.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p5en.48xlarge for training warm pool usage |
| ml.p6-b200.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for cluster spot instance usage |
| ml.p6-b200.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for cluster usage |
| ml.p6-b200.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for endpoint usage |
| ml.p6-b200.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.p6-b200.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for spot training job usage |
| ml.p6-b200.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for training job usage |
| ml.p6-b200.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p6-b200.48xlarge for training warm pool usage |
| ml.p6-b300.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p6-b300.48xlarge for cluster spot instance usage |
| ml.p6-b300.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p6-b300.48xlarge for cluster usage |
| ml.p6-b300.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.p6-b300.48xlarge for endpoint usage |
| ml.p6-b300.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.p6-b300.48xlarge for spot training job usage |
| ml.p6-b300.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.p6-b300.48xlarge for training job usage |
| ml.p6-b300.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p6-b300.48xlarge for training warm pool usage |
| ml.p6e-gb200.36xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.p6e-gb200.36xlarge for cluster spot instance usage |
| ml.p6e-gb200.36xlarge for cluster usage | Each supported Region: 0 | Yes | ml.p6e-gb200.36xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.p6e-gb200.36xlarge for training job usage | Each supported Region: 0 | Yes | ml.p6e-gb200.36xlarge for training job usage |
| ml.p6e-gb200.36xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.p6e-gb200.36xlarge for training warm pool usage |
| ml.r5.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5.12xlarge for endpoint usage |
| ml.r5.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r5.12xlarge for notebook instance usage |
| ml.r5.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.12xlarge for processing job usage |
| ml.r5.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.12xlarge for spot training job usage |
| ml.r5.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.12xlarge for training job usage |
| ml.r5.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.12xlarge for training warm pool usage |
| ml.r5.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r5.16xlarge for notebook instance usage |
| ml.r5.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.16xlarge for processing job usage |
| ml.r5.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.16xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.16xlarge for training job usage |
| ml.r5.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.16xlarge for training warm pool usage |
| ml.r5.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5.24xlarge for endpoint usage |
| ml.r5.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r5.24xlarge for notebook instance usage |
| ml.r5.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.24xlarge for processing job usage |
| ml.r5.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.24xlarge for spot training job usage |
| ml.r5.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.24xlarge for training job usage |
| ml.r5.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.24xlarge for training warm pool usage |
| ml.r5.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5.2xlarge for endpoint usage |
| ml.r5.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r5.2xlarge for notebook instance usage |
| ml.r5.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.2xlarge for processing job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.2xlarge for spot training job usage |
| ml.r5.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.2xlarge for training job usage |
| ml.r5.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.2xlarge for training warm pool usage |
| ml.r5.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5.4xlarge for endpoint usage |
| ml.r5.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r5.4xlarge for notebook instance usage |
| ml.r5.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.4xlarge for processing job usage |
| ml.r5.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.4xlarge for spot training job usage |
| ml.r5.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.4xlarge for training job usage |
| ml.r5.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.4xlarge for training warm pool usage |
| ml.r5.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r5.8xlarge for notebook instance usage |
| ml.r5.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.8xlarge for processing job usage |
| ml.r5.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.8xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.8xlarge for training job usage |
| ml.r5.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.8xlarge for training warm pool usage |
| ml.r5.large for endpoint usage | Each supported Region: 0 | Yes | ml.r5.large for endpoint usage |
| ml.r5.large for notebook instance usage | Each supported Region: 1 | Yes | ml.r5.large for notebook instance usage |
| ml.r5.large for processing job usage | Each supported Region: 0 | Yes | ml.r5.large for processing job usage |
| ml.r5.large for spot training job usage | Each supported Region: 0 | Yes | ml.r5.large for spot training job usage |
| ml.r5.large for training job usage | Each supported Region: 0 | Yes | ml.r5.large for training job usage |
| ml.r5.large for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.large for training warm pool usage |
| ml.r5.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5.xlarge for endpoint usage |
| ml.r5.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.r5.xlarge for notebook instance usage |
| ml.r5.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5.xlarge for processing job usage |
| ml.r5.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5.xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5.xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5.xlarge for training job usage |
| ml.r5.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5.xlarge for training warm pool usage |
| ml.r5d.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5d.12xlarge for endpoint usage |
| ml.r5d.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.12xlarge for processing job usage |
| ml.r5d.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.12xlarge for spot training job usage |
| ml.r5d.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.12xlarge for training job usage |
| ml.r5d.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.12xlarge for training warm pool usage |
| ml.r5d.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r5d.16xlarge for cluster spot instance usage |
| ml.r5d.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r5d.16xlarge for cluster usage |
| ml.r5d.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.16xlarge for processing job usage |
| ml.r5d.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.16xlarge for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5d.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.16xlarge for training job usage |
| ml.r5d.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.16xlarge for training warm pool usage |
| ml.r5d.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5d.24xlarge for endpoint usage |
| ml.r5d.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.24xlarge for processing job usage |
| ml.r5d.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.24xlarge for spot training job usage |
| ml.r5d.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.24xlarge for training job usage |
| ml.r5d.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.24xlarge for training warm pool usage |
| ml.r5d.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5d.2xlarge for endpoint usage |
| ml.r5d.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.2xlarge for processing job usage |
| ml.r5d.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.2xlarge for spot training job usage |
| ml.r5d.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.2xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5d.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.2xlarge for training warm pool usage |
| ml.r5d.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5d.4xlarge for endpoint usage |
| ml.r5d.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.4xlarge for processing job usage |
| ml.r5d.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.4xlarge for spot training job usage |
| ml.r5d.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.4xlarge for training job usage |
| ml.r5d.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.4xlarge for training warm pool usage |
| ml.r5d.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.8xlarge for processing job usage |
| ml.r5d.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.8xlarge for spot training job usage |
| ml.r5d.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.8xlarge for training job usage |
| ml.r5d.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.8xlarge for training warm pool usage |
| ml.r5d.large for endpoint usage | Each supported Region: 0 | Yes | ml.r5d.large for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r5d.large for processing job usage | Each supported Region: 0 | Yes | ml.r5d.large for processing job usage |
| ml.r5d.large for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.large for spot training job usage |
| ml.r5d.large for training job usage | Each supported Region: 0 | Yes | ml.r5d.large for training job usage |
| ml.r5d.large for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.large for training warm pool usage |
| ml.r5d.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r5d.xlarge for endpoint usage |
| ml.r5d.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r5d.xlarge for processing job usage |
| ml.r5d.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r5d.xlarge for spot training job usage |
| ml.r5d.xlarge for training job usage | Each supported Region: 0 | Yes | ml.r5d.xlarge for training job usage |
| ml.r5d.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r5d.xlarge for training warm pool usage |
| ml.r6g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.12xlarge for endpoint usage |
| ml.r6g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.16xlarge for endpoint usage |
| ml.r6g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.2xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r6g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.4xlarge for endpoint usage |
| ml.r6g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.8xlarge for endpoint usage |
| ml.r6g.large for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.large for endpoint usage |
| ml.r6g.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6g.xlarge for endpoint usage |
| ml.r6gd.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.12xlarge for endpoint usage |
| ml.r6gd.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.16xlarge for endpoint usage |
| ml.r6gd.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.2xlarge for endpoint usage |
| ml.r6gd.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.4xlarge for endpoint usage |
| ml.r6gd.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.8xlarge for endpoint usage |
| ml.r6gd.large for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.large for endpoint usage |
| ml.r6gd.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r6gd.xlarge for endpoint usage |
| ml.r6i.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.12xlarge for cluster spot instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r6i.12xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.12xlarge for cluster usage |
| ml.r6i.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.12xlarge for notebook instance usage |
| ml.r6i.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.12xlarge for transform job usage |
| ml.r6i.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.16xlarge for cluster spot instance usage |
| ml.r6i.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.16xlarge for cluster usage |
| ml.r6i.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.16xlarge for notebook instance usage |
| ml.r6i.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.16xlarge for transform job usage |
| ml.r6i.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.24xlarge for cluster spot instance usage |
| ml.r6i.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.24xlarge for cluster usage |
| ml.r6i.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.24xlarge for notebook instance usage |
| ml.r6i.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.24xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r6i.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.2xlarge for cluster spot instance usage |
| ml.r6i.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.2xlarge for cluster usage |
| ml.r6i.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.2xlarge for notebook instance usage |
| ml.r6i.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.2xlarge for transform job usage |
| ml.r6i.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.32xlarge for cluster spot instance usage |
| ml.r6i.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.32xlarge for cluster usage |
| ml.r6i.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.32xlarge for notebook instance usage |
| ml.r6i.32xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.32xlarge for transform job usage |
| ml.r6i.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.4xlarge for cluster spot instance usage |
| ml.r6i.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.4xlarge for cluster usage |
| ml.r6i.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.4xlarge for notebook instance usage |
| ml.r6i.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.4xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r6i.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.8xlarge for cluster spot instance usage |
| ml.r6i.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.8xlarge for cluster usage |
| ml.r6i.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6i.8xlarge for notebook instance usage |
| ml.r6i.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.8xlarge for transform job usage |
| ml.r6i.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.large for cluster spot instance usage |
| ml.r6i.large for cluster usage | Each supported Region: 0 | Yes | ml.r6i.large for cluster usage |
| ml.r6i.large for notebook instance usage | Each supported Region: 1 | Yes | ml.r6i.large for notebook instance usage |
| ml.r6i.large for transform job usage | Each supported Region: 0 | Yes | ml.r6i.large for transform job usage |
| ml.r6i.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r6i.xlarge for cluster spot instance usage |
| ml.r6i.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r6i.xlarge for cluster usage |
| ml.r6i.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.r6i.xlarge for notebook instance usage |
| ml.r6i.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r6i.xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r6id.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.12xlarge for notebook instance usage |
| ml.r6id.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.16xlarge for notebook instance usage |
| ml.r6id.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.24xlarge for notebook instance usage |
| ml.r6id.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.2xlarge for notebook instance usage |
| ml.r6id.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.32xlarge for notebook instance usage |
| ml.r6id.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.4xlarge for notebook instance usage |
| ml.r6id.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r6id.8xlarge for notebook instance usage |
| ml.r6id.large for notebook instance usage | Each supported Region: 1 | Yes | ml.r6id.large for notebook instance usage |
| ml.r6id.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.r6id.xlarge for notebook instance usage |
| ml.r7gd.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.12xlarge for endpoint usage |
| ml.r7gd.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.16xlarge for endpoint usage |
| ml.r7gd.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.2xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7gd.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.4xlarge for endpoint usage |
| ml.r7gd.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.8xlarge for endpoint usage |
| ml.r7gd.large for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.large for endpoint usage |
| ml.r7gd.medium for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.medium for endpoint usage |
| ml.r7gd.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7gd.xlarge for endpoint usage |
| ml.r7i.12xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for cluster spot instance usage |
| ml.r7i.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for endpoint usage |
| ml.r7i.12xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for notebook instance usage |
| ml.r7i.12xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for processing job usage |
| ml.r7i.12xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for spot training job usage |
| ml.r7i.12xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.12xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for training warm pool usage |
| ml.r7i.12xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.12xlarge for transform job usage |
| ml.r7i.16xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for cluster spot instance usage |
| ml.r7i.16xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for cluster usage |
| ml.r7i.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for endpoint usage |
| ml.r7i.16xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for notebook instance usage |
| ml.r7i.16xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for processing job usage |
| ml.r7i.16xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for spot training job usage |
| ml.r7i.16xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for training job usage |
| ml.r7i.16xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for training warm pool usage |
| ml.r7i.16xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.16xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.24xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for cluster spot instance usage |
| ml.r7i.24xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for cluster usage |
| ml.r7i.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for endpoint usage |
| ml.r7i.24xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for notebook instance usage |
| ml.r7i.24xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for processing job usage |
| ml.r7i.24xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for spot training job usage |
| ml.r7i.24xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for training job usage |
| ml.r7i.24xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for training warm pool usage |
| ml.r7i.24xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.24xlarge for transform job usage |
| ml.r7i.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for cluster spot instance usage |
| ml.r7i.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for cluster usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for endpoint usage |
| ml.r7i.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for notebook instance usage |
| ml.r7i.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for processing job usage |
| ml.r7i.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for spot training job usage |
| ml.r7i.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for training job usage |
| ml.r7i.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for training warm pool usage |
| ml.r7i.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.2xlarge for transform job usage |
| ml.r7i.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for cluster spot instance usage |
| ml.r7i.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for cluster usage |
| ml.r7i.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for endpoint usage |
| ml.r7i.48xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.48xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for processing job usage |
| ml.r7i.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for spot training job usage |
| ml.r7i.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for training job usage |
| ml.r7i.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for training warm pool usage |
| ml.r7i.48xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.48xlarge for transform job usage |
| ml.r7i.4xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for cluster spot instance usage |
| ml.r7i.4xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for cluster usage |
| ml.r7i.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for endpoint usage |
| ml.r7i.4xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for notebook instance usage |
| ml.r7i.4xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for processing job usage |
| ml.r7i.4xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for spot training job usage |
| ml.r7i.4xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.4xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for training warm pool usage |
| ml.r7i.4xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.4xlarge for transform job usage |
| ml.r7i.8xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for cluster spot instance usage |
| ml.r7i.8xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for cluster usage |
| ml.r7i.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for endpoint usage |
| ml.r7i.8xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for notebook instance usage |
| ml.r7i.8xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for processing job usage |
| ml.r7i.8xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for spot training job usage |
| ml.r7i.8xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for training job usage |
| ml.r7i.8xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for training warm pool usage |
| ml.r7i.8xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.8xlarge for transform job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.large for cluster spot instance usage |
| ml.r7i.large for cluster usage | Each supported Region: 0 | Yes | ml.r7i.large for cluster usage |
| ml.r7i.large for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.large for endpoint usage |
| ml.r7i.large for notebook instance usage | Each supported Region: 1 | Yes | ml.r7i.large for notebook instance usage |
| ml.r7i.large for processing job usage | Each supported Region: 0 | Yes | ml.r7i.large for processing job usage |
| ml.r7i.large for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.large for spot training job usage |
| ml.r7i.large for training job usage | Each supported Region: 0 | Yes | ml.r7i.large for training job usage |
| ml.r7i.large for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.large for training warm pool usage |
| ml.r7i.large for transform job usage | Each supported Region: 0 | Yes | ml.r7i.large for transform job usage |
| ml.r7i.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for cluster spot instance usage |
| ml.r7i.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for cluster usage |
| ml.r7i.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r7i.xlarge for notebook instance usage | Each supported Region: 1 | Yes | ml.r7i.xlarge for notebook instance usage |
| ml.r7i.xlarge for processing job usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for processing job usage |
| ml.r7i.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for spot training job usage |
| ml.r7i.xlarge for training job usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for training job usage |
| ml.r7i.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for training warm pool usage |
| ml.r7i.xlarge for transform job usage | Each supported Region: 0 | Yes | ml.r7i.xlarge for transform job usage |
| ml.r8g.12xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.12xlarge for endpoint usage |
| ml.r8g.16xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.16xlarge for endpoint usage |
| ml.r8g.24xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.24xlarge for endpoint usage |
| ml.r8g.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.2xlarge for endpoint usage |
| ml.r8g.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.48xlarge for endpoint usage |
| ml.r8g.4xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.4xlarge for endpoint usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.r8g.8xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.8xlarge for endpoint usage |
| ml.r8g.large for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.large for endpoint usage |
| ml.r8g.medium for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.medium for endpoint usage |
| ml.r8g.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.r8g.xlarge for endpoint usage |
| ml.t2.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.t2.2xlarge for endpoint usage |
| ml.t2.2xlarge for notebook instance usage | Each supported Region: 6 | Yes | ml.t2.2xlarge for notebook instance usage |
| ml.t2.large for endpoint usage | Each supported Region: 0 | Yes | ml.t2.large for endpoint usage |
| ml.t2.large for notebook instance usage | Each supported Region: 6 | Yes | ml.t2.large for notebook instance usage |
| ml.t2.medium for endpoint usage | Each supported Region: 2 | Yes | ml.t2.medium for endpoint usage |
| ml.t2.medium for notebook instance usage | Each supported Region: 6 | Yes | ml.t2.medium for notebook instance usage |
| ml.t2.xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.t2.xlarge for endpoint usage |
| ml.t2.xlarge for notebook instance usage | Each supported Region: 6 | Yes | ml.t2.xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.t3.2xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.t3.2xlarge for cluster spot instance usage |
| ml.t3.2xlarge for cluster usage | Each supported Region: 0 | Yes | ml.t3.2xlarge for cluster usage |
| ml.t3.2xlarge for notebook instance usage | Each supported Region: 6 | Yes | ml.t3.2xlarge for notebook instance usage |
| ml.t3.2xlarge for processing job usage | Each supported Region: 0 | Yes | ml.t3.2xlarge for processing job usage |
| ml.t3.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.t3.2xlarge for spot training job usage |
| ml.t3.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.t3.2xlarge for training job usage |
| ml.t3.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.t3.2xlarge for training warm pool usage |
| ml.t3.large for cluster spot instance usage | Each supported Region: 0 | Yes | ml.t3.large for cluster spot instance usage |
| ml.t3.large for cluster usage | Each supported Region: 0 | Yes | ml.t3.large for cluster usage |
| ml.t3.large for notebook instance usage | Each supported Region: 6 | Yes | ml.t3.large for notebook instance usage |
| ml.t3.large for processing job usage | Each supported Region: 4 | Yes | ml.t3.large for processing job usage |
| ml.t3.large for spot training job usage | Each supported Region: 0 | Yes | ml.t3.large for spot training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.t3.large for training job usage | Each supported Region: 0 | Yes | ml.t3.large for training job usage |
| ml.t3.large for training warm pool usage | Each supported Region: 0 | Yes | ml.t3.large for training warm pool usage |
| ml.t3.medium for cluster spot instance usage | Each supported Region: 0 | Yes | ml.t3.medium for cluster spot instance usage |
| ml.t3.medium for cluster usage | Each supported Region: 0 | Yes | ml.t3.medium for cluster usage |
| ml.t3.medium for notebook instance usage | Each supported Region: 6 | Yes | ml.t3.medium for notebook instance usage |
| ml.t3.medium for processing job usage | Each supported Region: 4 | Yes | ml.t3.medium for processing job usage |
| ml.t3.xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.t3.xlarge for cluster spot instance usage |
| ml.t3.xlarge for cluster usage | Each supported Region: 0 | Yes | ml.t3.xlarge for cluster usage |
| ml.t3.xlarge for notebook instance usage | Each supported Region: 6 | Yes | ml.t3.xlarge for notebook instance usage |
| ml.t3.xlarge for processing job usage | Each supported Region: 2 | Yes | ml.t3.xlarge for processing job usage |
| ml.t3.xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.t3.xlarge for spot training job usage |
| ml.t3.xlarge for training job usage | Each supported Region: 0 | Yes | ml.t3.xlarge for training job usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.t3.xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.t3.xlarge for training warm pool usage |
| ml.trn1.2xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.trn1.2xlarge for endpoint usage |
| ml.trn1.2xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.trn1.2xlarge for notebook instance usage |
| ml.trn1.2xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.trn1.2xlarge for spot training job usage |
| ml.trn1.2xlarge for training job usage | Each supported Region: 0 | Yes | ml.trn1.2xlarge for training job usage |
| ml.trn1.2xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.trn1.2xlarge for training warm pool usage |
| ml.trn1.2xlarge for transform job usage | Each supported Region: 0 | Yes | ml.trn1.2xlarge for transform job usage |
| ml.trn1.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for cluster spot instance usage |
| ml.trn1.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for cluster usage |
| ml.trn1.32xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for endpoint usage |
| ml.trn1.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for notebook instance usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.trn1.32xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for spot training job usage |
| ml.trn1.32xlarge for training job usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for training job usage |
| ml.trn1.32xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for training warm pool usage |
| ml.trn1.32xlarge for transform job usage | Each supported Region: 0 | Yes | ml.trn1.32xlarge for transform job usage |
| ml.trn1n.32xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for cluster spot instance usage |
| ml.trn1n.32xlarge for cluster usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for cluster usage |
| ml.trn1n.32xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for endpoint usage |
| ml.trn1n.32xlarge for notebook instance usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for notebook instance usage |
| ml.trn1n.32xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for spot training job usage |
| ml.trn1n.32xlarge for training job usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for training job usage |
| ml.trn1n.32xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.trn1n.32xlarge for training warm pool usage |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ml.trn2.3xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.trn2.3xlarge for cluster spot instance usage |
| ml.trn2.3xlarge for cluster usage | Each supported Region: 0 | Yes | ml.trn2.3xlarge for cluster usage |
| ml.trn2.48xlarge for cluster spot instance usage | Each supported Region: 0 | Yes | ml.trn2.48xlarge for cluster spot instance usage |
| ml.trn2.48xlarge for cluster usage | Each supported Region: 0 | Yes | ml.trn2.48xlarge for cluster usage |
| ml.trn2.48xlarge for endpoint usage | Each supported Region: 0 | Yes | ml.trn2.48xlarge for endpoint usage |
| ml.trn2.48xlarge for spot training job usage | Each supported Region: 0 | Yes | ml.trn2.48xlarge for spot training job usage |
| ml.trn2.48xlarge for training job usage | Each supported Region: 0 | Yes | ml.trn2.48xlarge for training job usage |
| ml.trn2.48xlarge for training warm pool usage | Each supported Region: 0 | Yes | ml.trn2.48xlarge for training warm pool usage |

# AWS Secrets Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | secretsmanager.us-east-2.amazonaws.com | HTTPS |
| | | secretsmanager-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | secretsmanager.us-east-1.amazonaws.com | HTTPS |
| | | secretsmanager-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | secretsmanager.us-west-1.amazonaws.com | HTTPS |
| | | secretsmanager-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | secretsmanager.us-west-2.amazonaws.com | HTTPS |
| | | secretsmanager-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | secretsmanager.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | secretsmanager.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | secretsmanager.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | secretsmanager.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | secretsmanager.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | secretsmanager.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | secretsmanager.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | secretsmanager.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | secretsmanager.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | secretsmanager.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | secretsmanager.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | secretsmanager.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | secretsmanager.ap-east-2.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | secretsmanager.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | secretsmanager.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | secretsmanager.ca-central-1.amazonaws.com<br>secretsmanager-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | secretsmanager.ca-west-1.amazonaws.com<br>secretsmanager-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | secretsmanager.eu-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | secretsmanager.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | secretsmanager.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | secretsmanager.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | secretsmanager.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | secretsmanager.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | secretsmanager.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | secretsmanager.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | secretsmanager.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | secretsmanager.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | secretsmanager.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | secretsmanager.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | secretsmanager.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | secretsmanager.us-gov-east-1.amazonaws.com<br><br>secretsmanager-fips.us-gov-east-1.amazonaws.com | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | secretsmanager.us-gov-west-1.amazonaws.com<br><br>secretsmanager-fips.us-gov-west-1.amazonaws.com | HTTPS<br><br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Combined rate of DeleteResourcePolicy, GetResourcePolicy, PutResourcePolicy, and ValidateResourcePolicy API requests | Each supported Region: 50 per second | No | The maximum transactions per second for DeleteResourcePolicy, GetResourcePolicy, PutResourcePolicy, and ValidateResourcePolicy API requests combined. |
| Combined rate of PutSecretValue, RemoveRegionsFromReplication, ReplicateSecretToRegion, StopRepli | Each supported Region: 50 per second | No | The maximum transactions per second for PutSecretValue, RemoveRegionsFromR |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| cationToReplica, UpdateSecret, and UpdateSecretVersionStage API requests | | | eplication, ReplicateSecretToRegion, StopReplicationToReplica, UpdateSecret, and UpdateSecretVersionStage API requests combined. |
| Combined rate of RestoreSecret API requests | Each supported Region: 50 per second | No | The maximum transactions per second for RestoreSecret API requests. |
| Combined rate of RotateSecret and CancelRotateSecret API requests | Each supported Region: 50 per second | No | The maximum transactions per second for RotateSecret and CancelRotateSecret API requests combined. |
| Combined rate of TagResource and UntagResource API requests | Each supported Region: 50 per second | No | The maximum transactions per second for TagResource and UntagResource API requests combined. |
| Rate of BatchGetSecretValue API requests | Each supported Region: 100 per second | No | The maximum transactions per second for BatchGetSecretValue API requests. |
| Rate of CreateSecret API requests | Each supported Region: 50 per second | No | The maximum transactions per second for CreateSecret API requests. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of DeleteSecret API requests | Each supported Region: 50 per second | No | The maximum transactions per second for DeleteSecret API requests. |
| Rate of DescribeSecret API requests | Each supported Region: 40,000 per second | No | The maximum transactions per second for DescribeSecret API requests. |
| Rate of GetRandomPassword API requests | Each supported Region: 50 per second | No | The maximum transactions per second for GetRandomPassword API requests. |
| Rate of GetSecretValue API requests | Each supported Region: 10,000 per second | No | The maximum transactions per second for GetSecretValue API requests. |
| Rate of ListSecretVersionIds API requests | Each supported Region: 50 per second | No | The maximum transactions per second for ListSecretVersionIds API requests. |
| Rate of ListSecrets API requests | Each supported Region: 100 per second | No | The maximum transactions per second for ListSecrets API requests. |
| Resource-based policy length | Each supported Region: 20,480 | No | The maximum number of characters in a resource-based permissions policy attached to a secret. |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Secret value size | Each supported Region: 65,536 Bytes | No | The maximum size of an encrypted secret value. If the secret value is a string, then this is the number of characters permitted in the secret value. |
| Secrets | Each supported Region: 500,000 | No | The maximum number of secrets in each AWS Region of this AWS account. |
| Staging labels attached across all versions of a secret | Each supported Region: 20 | No | The maximum number of staging labels attached across all versions of a secret. |
| Versions per secret | Each supported Region: 100 | No | The maximum number of versions of a secret. |

# Amazon Security Lake endpoints

The following are the service endpoints for this service. To connect programmatically to an AWS service, you use an endpoint. In addition to the standard AWS endpoints, some AWS services offer FIPS endpoints in selected Regions. For more information, see AWS service endpoints.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | securitylake.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | securitylake-fips.us-east-2.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-east-2.api.aws | HTTPS | |
| | | securitylake.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | securitylake.us-east-1.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-east-1.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-east-1.api.aws | HTTPS | |
| | | securitylake.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | securitylake.us-west-1.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-west-1.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-west-1.api.aws | HTTPS | |
| | | securitylake.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | securitylake.us-west-2.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-west-2.amazonaws.com | HTTPS | |
| | | securitylake-fips.us-west-2.api.aws | HTTPS | |
| | | securitylake.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | securitylake.ap-south-1.amazonaws.com | HTTPS | |
| | | securitylake.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | securitylake.ap-northeast-3.amazonaws.com | HTTPS | |
| | | securitylake.ap-northeast-3.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | securitylake.ap-northeast-2.amazonaws.com | HTTPS |
| | | securitylake.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | securitylake.ap-southeast-1.amazonaws.com | HTTPS |
| | | securitylake.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | securitylake.ap-southeast-2.amazonaws.com | HTTPS |
| | | securitylake.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | securitylake.ap-northeast-1.amazonaws.com | HTTPS |
| | | securitylake.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | securitylake.ca-central-1.amazonaws.com | HTTPS |
| | | securitylake.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | securitylake.eu-central-1.amazonaws.com | HTTPS |
| | | securitylake.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | securitylake.eu-west-1.amazonaws.com | HTTPS |
| | | securitylake.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | securitylake.eu-west-2.amazonaws.com | HTTPS |
| | | securitylake.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | securitylake.eu-west-3.amazonaws.com | HTTPS |
| | | securitylake.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | securitylake.eu-north-1.amazonaws.com<br>securitylake.eu-north-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | securitylake.sa-east-1.amazonaws.com<br>securitylake.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | securitylake.us-gov-east-1.amazonaws.com<br>securitylake.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | securitylake.us-gov-west-1.amazonaws.com<br>securitylake.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

# AWS Security Hub CSPM endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | securityhub.us-east-2.amazonaws.com | HTTPS |
| | | securityhub-fips.us-east-2.amazonaws.com | HTTPS |
| | | securityhub.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | securityhub.us-east-1.amazonaws.com | HTTPS |
| | | securityhub-fips.us-east-1.amazonaws.com | HTTPS |
| | | securityhub.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | securityhub.us-west-1.amazonaws.com | HTTPS |
| | | securityhub-fips.us-west-1.amazonaws.com | HTTPS |
| | | securityhub.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | securityhub.us-west-2.amazonaws.com | HTTPS |
| | | securityhub-fips.us-west-2.amazonaws.com | HTTPS |
| | | securityhub.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | securityhub.af-south-1.amazonaws.com | HTTPS |
| | | securityhub.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | securityhub.ap-east-1.amazonaws.com | HTTPS |
| | | securityhub.ap-east-1.api.aws | HTTPS |
| Asia Pacific | ap-south-2 | securityhub.ap-south-2.amazonaws.com | HTTPS |
| | | securityhub.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | securityhub.ap-southeast-3.amazonaws.com<br>securityhub.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | securityhub.ap-southeast-5.amazonaws.com<br>securityhub.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | securityhub.ap-southeast-4.amazonaws.com<br>securityhub.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | securityhub.ap-south-1.amazonaws.com<br>securityhub.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | securityhub.ap-southeast-6.amazonaws.com<br>securityhub.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | securityhub.ap-northeast-3.amazonaws.com<br>securityhub.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | securityhub.ap-northeast-2.amazonaws.com<br>securityhub.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | securityhub.ap-southeast-1.amazonaws.com<br>securityhub.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | securityhub.ap-southeast-2.amazonaws.com<br>securityhub.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | securityhub.ap-east-2.amazonaws.com<br>securityhub.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | securityhub.ap-southeast-7.amazonaws.com<br>securityhub.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | securityhub.ap-northeast-1.amazonaws.com<br>securityhub.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | securityhub.ca-central-1.amazonaws.com<br>securityhub.ca-central-1.api.aws | HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | securityhub.ca-west-1.amazonaws.com<br>securityhub.ca-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | securityhub.eu-central-1.amazonaws.com<br>securityhub.eu-central-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | securityhub.eu-west-1.amazonaws.com | HTTPS | |
| | | securityhub.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | securityhub.eu-west-2.amazonaws.com | HTTPS | |
| | | securityhub.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | securityhub.eu-south-1.amazonaws.com | HTTPS | |
| | | securityhub.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | securityhub.eu-west-3.amazonaws.com | HTTPS | |
| | | securityhub.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | securityhub.eu-south-2.amazonaws.com | HTTPS | |
| | | securityhub.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | securityhub.eu-north-1.amazonaws.com | HTTPS | |
| | | securityhub.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | securityhub.eu-central-2.amazonaws.com | HTTPS | |
| | | securityhub.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | securityhub.il-central-1.amazonaws.com | HTTPS | |
| | | securityhub.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | securityhub.mx-central-1.amazonaws.com | HTTPS | |
| | | securityhub.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | securityhub.me-south-1.amazonaws.com<br>securityhub.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | securityhub.me-central-1.amazonaws.com<br>securityhub.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | securityhub.sa-east-1.amazonaws.com<br>securityhub.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | securityhub.us-gov-east-1.amazonaws.com<br>securityhub-fips.us-gov-east-1.amazonaws.com<br>securityhub.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | securityhub.us-gov-west-1.amazonaws.com<br>securityhub-fips.us-gov-west-1.amazonaws.com<br>securityhub.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of Security Hub member accounts | Each supported Region: 10,000 | No | The maximum number of Security Hub member accounts that can be |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | added for each Security Hub administrator account in each Region. |
| Number of Security Hub outstanding invitations | Each supported Region: 1,000 | No | The maximum number of outstanding Security Hub member account invitations that can be sent per AWS account (Security Hub administrator account) per Region. |
| Number of automation rules | Each supported Region: 100 | No | The maximum number of automation rules that can be created by a Security Hub administrator account. |
| Number of custom actions | Each supported Region: 50 | No | The maximum number of custom actions that can be created per account per Region. |
| Number of custom insights | Each supported Region: 100 | No | The maximum number of user-defined custom insights that can be created per AWS account per Region. |
| Number of insight results | Each supported Region: 100 | No | The maximum number of aggregated results returned for the GetInsightsResults API operation. |

| Name | Default | Adjuste | Description |
| --- | --- | --- | --- |
| Security Hub finding retention time | Each supported Region: 90 | No | The maximum number of days a Security Hub finding is saved. This is 90 days after the most recent update or 90 days after the creation date if no update occurs. |

For more information about Security Hub CSPM quotas, see Quotas in the *AWS Security Hub User Guide*.

# AWS Security Incident Response endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | security-ir.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | security-ir.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (Oregon) | us-west-2 | security-ir.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | security-ir.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | security-ir.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | security-ir.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | security-ir.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | security-ir.ap-southeast-4.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | security-ir.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | security-ir.ap-northeast-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | security-ir.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | security-ir.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | security-ir.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | security-ir.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | security-ir.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | security-ir.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | security-ir.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | security-ir.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | security-ir.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | security-ir.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | security-ir.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | security-ir.eu-central-2.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | security-ir.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | security-ir.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | security-ir.sa-east-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Active self managed cases | Each supported Region: 100 | Yes | The number of active self managed cases. |
| Active service managed cases | Each supported Region: 50 | Yes | The number of active service managed cases. |
| Service managed cases created in the last day | Each supported Region: 50 | Yes | The number of service managed cases created in the last day. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate limit for BatchGetMemberAccountDetails | Each supported Region: 5 per second | No | The maximum number of BatchGetMemberAccountDetails requests you can make per second per account, in the current region |
| Throttle rate limit for CancelMembership | Each supported Region: 1 per second | No | The maximum number of CancelMembership requests you can make per second per account, in the current region |
| Throttle rate limit for CloseCase | Each supported Region: 1 per second | No | The maximum number of CloseCase requests you can make per second per account, in the current region |
| Throttle rate limit for CreateCase | Each supported Region: 1 per second | No | The maximum number of CreateCase requests you can make per second per account, in the current region |
| Throttle rate limit for CreateCaseComment | Each supported Region: 1 per second | No | The maximum number of CreateCaseComment requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate limit for CreateMembership | Each supported Region: 1 per second | No | The maximum number of CreateMembership requests you can make per second per account, in the current region |
| Throttle rate limit for GetCase | Each supported Region: 5 per second | No | The maximum number of GetCase requests you can make per second per account, in the current region |
| Throttle rate limit for GetCaseAttachmentDownloadUrl | Each supported Region: 1 per second | No | The maximum number of GetCaseAttachmentDownloadUrl requests you can make per second per account, in the current region |
| Throttle rate limit for GetCaseAttachmentUploadUrl | Each supported Region: 1 per second | No | The maximum number of GetCaseAttachmentUploadUrl requests you can make per second per account, in the current region |
| Throttle rate limit for GetMembership | Each supported Region: 5 per second | No | The maximum number of GetMembership requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for ListCaseEdits | Each supported Region: 5 per second | No | The maximum number of ListCaseEdits requests you can make per second per account, in the current region |
| Throttle rate limit for ListCases | Each supported Region: 5 per second | No | The maximum number of ListCases requests you can make per second per account, in the current region |
| Throttle rate limit for ListComments | Each supported Region: 5 per second | No | The maximum number of ListComments requests you can make per second per account, in the current region |
| Throttle rate limit for ListInvestigations | Each supported Region: 5 per second | No | The maximum number of ListInvestigations requests you can make per second per account, in the current region |
| Throttle rate limit for ListMemberships | Each supported Region: 5 per second | No | The maximum number of ListMemberships requests you can make per second per account, in the current region |
| Throttle rate limit for SendFeedback | Each supported Region: 1 per second | No | The maximum number of SendFeedback requests you can make per second per account, in the current region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate limit for UpdateCase | Each supported Region: 1 per second | No | The maximum number of UpdateCase requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateCaseComment | Each supported Region: 1 per second | No | The maximum number of UpdateCaseComment requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateCaseStatus | Each supported Region: 1 per second | No | The maximum number of UpdateCaseStatus requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateMembership | Each supported Region: 1 per second | No | The maximum number of UpdateMembership requests you can make per second per account, in the current region |
| Throttle rate limit for UpdateResolverType | Each supported Region: 1 per second | No | The maximum number of UpdateResolverType requests you can make per second per account, in the current region |

# Service Quotas endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | servicequotas.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | servicequotas.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | servicequotas.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | servicequotas.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | servicequotas.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | servicequotas.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific | ap-south-2 | servicequotas.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | servicequotas.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | servicequotas.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | servicequotas.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | servicequotas.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | servicequotas.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | servicequotas.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | servicequotas.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | servicequotas.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | servicequotas.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | servicequotas.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | servicequotas.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | servicequotas.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | servicequotas.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | servicequotas.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | servicequotas.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | servicequotas.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | servicequotas.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | servicequotas.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | servicequotas.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | servicequotas.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | servicequotas.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | servicequotas.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | servicequotas.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | servicequotas.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | servicequotas.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | servicequotas.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | servicequotas.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | servicequotas.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | servicequotas.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Active requests per account | Each supported Region: 20 | No | The maximum number of active service quota increase requests allowed per account |
| Active requests per account per opt-in Region | eu-south-2: 10<br><br>Each of the other supported Regions: 2 | No | The maximum number of active service quota increase requests allowed per account in opt-in Regions. For more information about opt-in Regions, see https://docs.aws.amazon.com/accounts/latest/r |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | eference/manage-acct-regions.html |
| Active requests per quota | Each supported Region: 1 | No | The maximum number of active service quota increase requests per quota, in the current Region |
| Max requests per template | Each supported Region: 10 | No | The maximum number of service quota increase requests allowed per template |
| Throttle Rate for AssociateServiceQuotaTemplate | Each supported Region: 1 per second | No | The maximum number of AssociateServiceQuotaTemplate requests allowed per second per account, in the current Region |
| Throttle rate for DeleteServiceQuotaIncreaseRequestFromTemplate | Each supported Region: 2 per second | No | The maximum number of DeleteServiceQuotaIncreaseRequestFromTemplate requests allowed per second per account, in the current Region |
| Throttle rate for DisassociateServiceQuotaTemplate | Each supported Region: 1 per second | No | The maximum number of DisassociateServiceQuotaTemplate requests allowed per second per account, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate for GetAWSDefaultServiceQuota | Each supported Region: 5 per second | No | The maximum number of GetAWSDefaultServiceQuota requests allowed per second per account, in the current Region |
| Throttle rate for GetAssociationForServiceQuotaTemplate | Each supported Region: 2 per second | No | The maximum number of GetAssociationForServiceQuotaTemplate requests allowed per second per account, in the current Region |
| Throttle rate for GetRequestedServiceQuotaChange | Each supported Region: 5 per second | No | The maximum number of GetRequestedServiceQuotaChange requests allowed per second per account, in the current Region |
| Throttle rate for GetServiceQuota | Each supported Region: 5 per second | No | The maximum number of GetServiceQuota requests allowed per second per account, in the current Region |
| Throttle rate for GetServiceQuotaIncreaseRequestFromTemplate | Each supported Region: 2 per second | No | The maximum number of GetServiceQuotaIncreaseRequestFromTemplate requests allowed per second per account, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate for ListAWSDefaultServiceQuotas | Each supported Region: 10 per second | No | The maximum number of ListAWSDefaultServiceQuotas requests allowed per second per account, in the current Region |
| Throttle rate for ListRequestedServiceQuotaChangeHistory | Each supported Region: 5 per second | No | The maximum number of ListRequestedServiceQuotaChangeHistory requests allowed per second per account, in the current Region |
| Throttle rate for ListRequestedServiceQuotaChangeHistoryByQuota | Each supported Region: 5 per second | No | The maximum number of ListRequestedServiceQuotaChangeHistoryByQuota requests allowed per second per account, in the current Region |
| Throttle rate for ListServiceQuotaIncreaseRequestsInTemplate | Each supported Region: 2 per second | No | The maximum number of ListServiceQuotaIncreaseRequestsInTemplate requests allowed per second per account, in the current Region |
| Throttle rate for ListServiceQuotas | Each supported Region: 10 per second | No | The maximum number of ListServiceQuotas requests allowed per second per account, in the current Region |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Throttle rate for ListServices | Each supported Region: 10 per second | No | The maximum number of ListServices requests allowed per second per account, in the current Region |
| Throttle rate for ListTagsForResource | Each supported Region: 10 per second | No | The maximum number of ListTagsForResource requests allowed per second per account, in the current Region |
| Throttle rate for PutServiceQuotaIncreaseRequestIntoTemplate | Each supported Region: 1 per second | No | The maximum number of PutServiceQuotaIncreaseRequestIntoTemplate requests allowed per second per account, in the current Region |
| Throttle rate for RequestServiceQuotaIncrease | Each supported Region: 3 per second | No | The maximum number of RequestServiceQuotaIncrease requests allowed per second per account, in the current Region |
| Throttle rate for TagResource | Each supported Region: 10 per second | No | The maximum number of TagResource requests allowed per second per account, in the current Region |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Throttle rate for UntagResource | Each supported Region: 10 per second | No | The maximum number of UntagResource requests allowed per second per account, in the current Region |

# AWS Serverless Application Repository endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | serverlessrepo.us-east-2.amazonaws.com | HTTPS |
| | | serverlessrepo-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | serverlessrepo.us-east-1.amazonaws.com | HTTPS |
| | | serverlessrepo-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | serverlessrepo.us-west-1.amazonaws.com | HTTPS |
| | | serverlessrepo-fips.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | serverlessrepo.us-west-2.amazonaws.com<br>serverlessrepo-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | serverlessrepo.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | serverlessrepo.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | serverlessrepo.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | serverlessrepo.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | serverlessrepo.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | serverlessrepo.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | serverlessrepo.ca-central-1.amazonaws.com<br>serverlessrepo-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | serverlessrepo.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | serverlessrepo.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | serverlessrepo.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | serverlessrepo.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | serverlessrepo.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | serverlessrepo.me-south-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | serverlessrepo.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | serverlessrepo.us-gov-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | serverlessrepo.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Application policy length | Each supported Region: 6,144 | No | The maximum application policy length (in characters). |
| Free Amazon S3 storage for code packages | Each supported Region: 5 Gigabytes | No | The maximum amount (in GB) of free Amazon S3 storage for code packages per AWS account per AWS region. |
| Public applications | Each supported Region: 100 | Yes | The maximum number of public applications per AWS account per AWS region. |

For more information, see AWS Serverless Application Repository Quotas in the *AWS Serverless Application Repository Developer Guide*.

# AWS Service Catalog endpoints and quotas

The following are the service endpoints and service quotas for this service. To connect programmatically to an AWS service, you use an endpoint. In addition to the standard AWS endpoints, some AWS services offer FIPS endpoints in selected Regions. For more information, see Amazon service endpoints. Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see Amazon service quotas.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | servicecatalog.us-east-2.amazonaws.com<br>servicecatalog-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | servicecatalog.us-east-1.amazonaws.com<br>servicecatalog-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | servicecatalog.us-west-1.amazonaws.com<br>servicecatalog-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | servicecatalog.us-west-2.amazonaws.com<br>servicecatalog-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | servicecatalog.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | servicecatalog.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | servicecatalog.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | servicecatalog.ap-southeast-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | servicecatalog.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | servicecatalog.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | servicecatalog.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | servicecatalog.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | servicecatalog.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | servicecatalog.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | servicecatalog.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | servicecatalog.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | servicecatalog.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | servicecatalog.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | servicecatalog.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | servicecatalog.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | servicecatalog.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | servicecatalog.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | servicecatalog.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | servicecatalog.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | servicecatalog.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | servicecatalog.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | servicecatalog.il-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | servicecatalog.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | servicecatalog.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | servicecatalog.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | servicecatalog.us-gov-east-1.amazonaws.com<br>servicecatalog-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | servicecatalog.us-gov-west-1.amazonaws.com<br>servicecatalog-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Applications per attribute group | Each supported Region: 1,000 | Yes | The maximum number of applications per attribute group |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Applications per region | Each supported Region: 2,000 | Yes | The maximum number of applications you can create per region |
| Attribute groups per application | Each supported Region: 1,000 | Yes | The maximum number of attribute groups per application |
| Attribute groups per region | Each supported Region: 2,000 | Yes | The maximum number of attribute groups you can create per region |
| Delegated administrators per organization | Each supported Region: 50 | No | The maximum number of delegated administrators you can register per organization |
| Portfolios per region | Each supported Region: 100 | Yes | The maximum number of portfolios you can create per region |
| Product versions per product | Each supported Region: 100 | Yes | The maximum number of product versions you can create per product |
| Products per portfolio | Each supported Region: 150 | Yes | The maximum number of products you can create per portfolio |
| Products per region | Each supported Region: 350 | Yes | The maximum number of products you can create per region |
| Resources per application | Each supported Region: 1,000 | Yes | The maximum number of resources per applications |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Service action associations per provisioning artifact | Each supported Region: 25 | No | The maximum number of service action associations you can create per provisioning artifact |
| Service actions per region | Each supported Region: 200 | No | The maximum number of service actions you can create per region |
| Shared accounts per portfolio | Each supported Region: 5,000 | No | The maximum number of shared accounts per portfolio |
| TagOptions per resource | Each supported Region: 25 | No | The maximum number of TagOptions you can associate with a resource |
| Tags per portfolio | Each supported Region: 20 | No | The maximum number of tags you can create per portfolio |
| Tags per product | Each supported Region: 20 | No | The maximum number of tags you can create per product |
| Tags per provisioned product | Each supported Region: 50 | No | The maximum number of tags you can create per provisioned product |
| Users, groups, and roles per portfolio | Each supported Region: 100 | Yes | The maximum number of users, groups, and roles you can create per portfolio |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Users, groups, and roles per product | Each supported Region: 200 | Yes | The maximum number of users, groups, and roles you can create per portfolio |
| Values per TagOption | Each supported Region: 25 | No | The maximum number of different values for each TagOption |

For more information, see Service Catalog default service quotas in the *Service Catalog Administrator Guide*.

# AWS Shield Advanced endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | shield.us-east-1.amazonaws.com | HTTPS | |
| | | shield-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Ireland) | eu-west-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (London) | eu-west-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Milan) | eu-south-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Paris) | eu-west-3 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Spain) | eu-south-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | shield.us-east-1.amazonaws.com<br>shield-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AWS Global Accelerator accelerator protections | Each supported Region: 1,000 | Yes | The maximum number of AWS Global Accelerator accelerators you can monitor and protect. |
| Amazon Route 53 hosted zone protections | Each supported Region: 1,000 | Yes | The maximum number of Amazon Route 53 hosted zones you can monitor and protect. |
| CloudFront distribution protections | Each supported Region: 1,000 | Yes | The maximum number of Amazon CloudFront distributions you can monitor and protect. |
| Elastic IP address protections | Each supported Region: 1,000 | Yes | The maximum number of Elastic IP addresses you can monitor and protect. |
| Elastic Load Balancing load balancer protections | Each supported Region: 1,000 | Yes | The maximum number of Elastic Load Balancing load balancers you can monitor and protect. |

# Amazon Simple Storage Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## Amazon S3 endpoints

When you use the REST API to send requests to the endpoints shown in the following table, you can use the virtual-hosted style and path-style methods. For more information, see Virtual hosting of buckets.

> **Note**
>
> Some Regions support legacy endpoints. For information about Amazon S3 backward compatibility for path-style and virtual-hosted–style URL requests, see Backward compatibility.

### Amazon S3 regular endpoints

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | s3.us-east-2.amazonaws.com<br>s3.dualstack.us-east-2.amazonaws.com<br>s3-fips.dualstack.us-east-2.amazonaws.com<br>s3-fips.us-east-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTPS | 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | s3.us-east-1.amazonaws.com<br>s3.dualstack.us-east-1.amazonaws.com<br>s3-fips.us-east-1.amazonaws.com<br>s3-fips.dualstack.us-east-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTPS | 2 & 4 |
| US West (N. California) | us-west-1 | s3.us-west-1.amazonaws.com<br>s3.dualstack.us-west-1.amazonaws.com<br>s3-fips.dualstack.us-west-1.amazonaws.com<br>s3-fips.us-west-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTPS | 2 & 4 |
| US West (Oregon) | us-west-2 | s3.us-west-2.amazonaws.com<br>s3.dualstack.us-west-2.amazonaws.com<br>s3-fips.dualstack.us-west-2.amazonaws.com<br>s3-fips.us-west-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTPS | 2 & 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | s3.af-south-1.amazonaws.com<br>s3.dualstack.af-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Hong Kong) | ap-east-1 | s3.ap-east-1.amazonaws.com<br>s3.dualstack.ap-east-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Hyderabad) | ap-south-2 | s3.ap-south-2.amazonaws.com<br>s3.dualstack.ap-south-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Jakarta) | ap-southeast-3 | s3.ap-southeast-3.amazonaws.com<br>s3.dualstack.ap-southeast-3.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | s3.ap-southeast-5.amazonaws.com<br>s3.dualstack.ap-southeast-5.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Melbourne) | ap-southeast-4 | s3.ap-southeast-4.amazonaws.com<br>s3.dualstack.ap-southeast-4.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Mumbai) | ap-south-1 | s3.ap-south-1.amazonaws.com<br>s3.dualstack.ap-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (New Zealand) | ap-southeast-6 | s3.ap-southeast-6.amazonaws.com<br>s3.dualstack.ap-southeast-6.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | s3.ap-northeast-3.amazonaws.com<br>s3.dualstack.ap-northeast-3.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Seoul) | ap-northeast-2 | s3.ap-northeast-2.amazonaws.com<br>s3.dualstack.ap-northeast-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Singapore) | ap-southeast-1 | s3.ap-southeast-1.amazonaws.com<br>s3.dualstack.ap-southeast-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 2 & 4 |
| Asia Pacific (Sydney) | ap-southeast-2 | s3.ap-southeast-2.amazonaws.com<br>s3.dualstack.ap-southeast-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 2 & 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | s3.ap-east-2.amazonaws.com<br>s3.dualstack.ap-east-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Thailand) | ap-southeast-7 | s3.ap-southeast-7.amazonaws.com<br>s3.dualstack.ap-southeast-7.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Asia Pacific (Tokyo) | ap-northeast-1 | s3.ap-northeast-1.amazonaws.com<br>s3.dualstack.ap-northeast-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 2 & 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | s3.ca-central-1.amazonaws.com<br>s3.dualstack.ca-central-1.amazonaws.com<br>s3-fips.dualstack.ca-central-1.amazonaws.com<br>s3-fips.ca-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTPS | 4 |
| Canada West (Calgary) | ca-west-1 | s3.ca-west-1.amazonaws.com<br>s3.dualstack.ca-west-1.amazonaws.com<br>s3-fips.dualstack.ca-west-1.amazonaws.com<br>s3-fips.ca-west-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS<br>HTTPS<br>HTTPS | 4 |
| Europe (Frankfurt) | eu-central-1 | s3.eu-central-1.amazonaws.com<br>s3.dualstack.eu-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | s3.eu-west-1.amazonaws.com<br>s3.dualstack.eu-west-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 2 & 4 |
| Europe (London) | eu-west-2 | s3.eu-west-2.amazonaws.com<br>s3.dualstack.eu-west-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Europe (Milan) | eu-south-1 | s3.eu-south-1.amazonaws.com<br>s3.dualstack.eu-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Europe (Paris) | eu-west-3 | s3.eu-west-3.amazonaws.com<br>s3.dualstack.eu-west-3.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | s3.eu-south-2.amazonaws.com<br>s3.dualstack.eu-south-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Europe (Stockholm) | eu-north-1 | s3.eu-north-1.amazonaws.com<br>s3.dualstack.eu-north-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Europe (Zurich) | eu-central-2 | s3.eu-central-2.amazonaws.com<br>s3.dualstack.eu-central-2.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Israel (Tel Aviv) | il-central-1 | s3.il-central-1.amazonaws.com<br>s3.dualstack.il-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| Mexico (Central) | mx-central-1 | s3.mx-central-1.amazonaws.com<br>s3.dualstack.mx-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Middle East (Bahrain) | me-south-1 | s3.me-south-1.amazonaws.com<br>s3.dualstack.me-south-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| Middle East (UAE) | me-central-1 | s3.me-central-1.amazonaws.com<br>s3.dualstack.me-central-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 4 |
| South America (São Paulo) | sa-east-1 | s3.sa-east-1.amazonaws.com<br>s3.dualstack.sa-east-1.amazonaws.com | HTTP and HTTPS<br>HTTP and HTTPS | 2 & 4 |

| Region Name | Region | Endpoint | Protocol | Signature Version(s) Support |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | s3.us-gov-east-1.amazonaws.com<br>s3-fips.dualstack.us-gov-east-1.amazonaws.com<br>s3.dualstack.us-gov-east-1.amazonaws.com<br>s3-fips.us-gov-east-1.amazonaws.com | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS<br>HTTP | 4 |
| AWS GovCloud (US-West) | us-gov-west-1 | s3.us-gov-west-1.amazonaws.com<br>s3-fips.dualstack.us-gov-west-1.amazonaws.com<br>s3.dualstack.us-gov-west-1.amazonaws.com<br>s3-fips.us-gov-west-1.amazonaws.com | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS<br>HTTP | 2 & 4 |

### Amazon S3 Tables endpoints

The `s3-tables` endpoints are used only with Amazon S3 Tables operations. For more information, see Working with Amazon S3 Tables and table buckets.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | s3tables.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | s3tables.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | s3tables.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | s3tables.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | s3tables.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | s3tables.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | s3tables.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | s3tables.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific | ap-southeast-5 | s3tables.ap-southeast-5.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Malaysia) | | | | |
| Asia Pacific (Melbourne) | ap-southeast-4 | s3tables.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | s3tables.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | s3tables.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | s3tables.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | s3tables.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | s3tables.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | s3tables.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | s3tables.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | s3tables.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | s3tables.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | s3tables.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | s3tables.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | s3tables.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | s3tables.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | s3tables.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | s3tables.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | s3tables.eu-west-3.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | s3tables.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | s3tables.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | s3tables.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | s3tables.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | s3tables.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | s3tables.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | s3tables.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | s3tables.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | s3tables.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | s3tables.us-gov-west-1.amazonaws.com | HTTPS | |

### Amazon S3 Access Point endpoints

The `s3-accesspoint` endpoints are used only to make requests through Amazon S3 Access Points. For more information, see Working with Amazon S3 Access Points.

To interact with these endpoints you must precede the endpoint name with the relevant Access Point name in one of the following formats:

*access-point*.s3-accesspoint.*AWS Region code*.amazonaws.com

*access-point*.s3-accesspoint-fips.dualstack.*AWS Region code*.amazonaws.com

*access-point*.s3-accesspoint-fips.*AWS Region code*.amazonaws.com

*access-point*.s3-accesspoint.dualstack.*AWS Region code*.amazonaws.com

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | s3-accesspoint.us-east-2.amazonaws.com<br>s3-accesspoint-fips.dualstack.us-east-2.amazonaws.com<br>s3-accesspoint-fips.us-east-2.amazonaws.com<br>s3-accesspoint.dualstack.us-east-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | s3-accesspoint.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | s3-accesspoint-fips.dualstack.us-east-1.amazonaws.com<br>s3-accesspoint-fips.us-east-1.amazonaws.com<br>s3-accesspoint.dualstack.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | s3-accesspoint.us-west-1.amazonaws.com<br>s3-accesspoint-fips.dualstack.us-west-1.amazonaws.com<br>s3-accesspoint-fips.us-west-1.amazonaws.com<br>s3-accesspoint.dualstack.us-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | s3-accesspoint.us-west-2.amazonaws.com<br>s3-accesspoint-fips.dualstack.us-west-2.amazonaws.com<br>s3-accesspoint-fips.us-west-2.amazonaws.com<br>s3-accesspoint.dualstack.us-west-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | s3-accesspoint.af-south-1.amazonaws.com<br>s3-accesspoint.dualstack.af-south-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | s3-accesspoint.ap-east-1.amazonaws.com<br>s3-accesspoint.dualstack.ap-east-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | s3-accesspoint.ap-south-2.amazonaws.com<br>s3-accesspoint.dualstack.ap-south-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | s3-accesspoint.ap-southeast-3.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-3.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | s3-accesspoint.ap-southeast-5.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-5.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | s3-accesspoint.ap-southeast-4.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-4.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | s3-accesspoint.ap-south-1.amazonaws.com<br>s3-accesspoint.dualstack.ap-south-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | s3-accesspoint.ap-southeast-6.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-6.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | s3-accesspoint.ap-northeast-3.amazonaws.com<br>s3-accesspoint.dualstack.ap-northeast-3.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | s3-accesspoint.ap-northeast-2.amazonaws.com<br>s3-accesspoint.dualstack.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | s3-accesspoint.ap-southeast-1.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | s3-accesspoint.ap-southeast-2.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | s3-accesspoint.ap-east-2.amazonaws.com<br>s3-accesspoint.dualstack.ap-east-2.amazonaws.com | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | s3-accesspoint.ap-southeast-7.amazonaws.com<br>s3-accesspoint.dualstack.ap-southeast-7.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | s3-accesspoint.ap-northeast-1.amazonaws.com<br>s3-accesspoint.dualstack.ap-northeast-1.amazonaws.com | HTTPS<br>HTTPS | |
| Canada (Central) | ca-central-1 | s3-accesspoint.ca-central-1.amazonaws.com<br>s3-accesspoint-fips.dualstack.ca-central-1.amazonaws.com<br>s3-accesspoint-fips.ca-central-1.amazonaws.com<br>s3-accesspoint.dualstack.ca-central-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Canada West (Calgary) | ca-west-1 | s3-accesspoint.ca-west-1.amazonaws.com<br>s3-accesspoint-fips.dualstack.ca-west-1.amazonaws.com<br>s3-accesspoint-fips.ca-west-1.amazonaws.com<br>s3-accesspoint.dualstack.ca-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | s3-accesspoint.eu-central-1.amazonaws.com<br>s3-accesspoint.dualstack.eu-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | s3-accesspoint.eu-west-1.amazonaws.com<br>s3-accesspoint.dualstack.eu-west-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | s3-accesspoint.eu-west-2.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | s3-accesspoint.eu-south-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | s3-accesspoint.eu-west-3.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | s3-accesspoint.eu-south-2.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | s3-accesspoint.eu-north-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | s3-accesspoint.eu-central-2.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | s3-accesspoint.il-central-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.il-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | s3-accesspoint.mx-central-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.mx-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | s3-accesspoint.me-south-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | s3-accesspoint.me-central-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | s3-accesspoint.sa-east-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | s3-accesspoint.us-gov-east-1.amazonaws.com | HTTPS |
| | | s3-accesspoint-fips.dualstack.us-gov-east-1.amazonaws.com | HTTPS |
| | | s3-accesspoint-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | s3-accesspoint.dualstack.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | s3-accesspoint.us-gov-west-1.amazonaws.com<br>s3-accesspoint-fips.dualstack.us-gov-west-1.amazonaws.com<br>s3-accesspoint-fips.us-gov-west-1.amazonaws.com<br>s3-accesspoint.dualstack.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

### Amazon S3 Control endpoints

The `s3-control` endpoints are used with Amazon S3 account-level operations. To interact with these endpoints you must precede the endpoint name with the relevant AWS account ID in one of the following formats:

*account-id*.s3-control.*AWS Region code*.amazonaws.com

*account-id*.s3-control.dualstack.*AWS Region code*.amazonaws.com

*account-id*.s3-control-fips.dualstack.*AWS Region code*.amazonaws.com

*account-id*.s3-control-fips.*AWS Region code*.amazonaws.com

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | s3-control.us-east-2.amazonaws.com<br>s3-control.dualstack.us-east-2.amazonaws.com<br>s3-control-fips.dualstack.us-east-2.amazonaws.com<br>s3-control-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | s3-control.us-east-1.amazonaws.com<br>s3-control-fips.us-east-1.amazonaws.com<br>s3-control.dualstack.us-east-1.amazonaws.com<br>s3-control-fips.dualstack.us-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (N. California) | us-west-1 | s3-control.us-west-1.amazonaws.com<br>s3-control.dualstack.us-west-1.amazonaws.com<br>s3-control-fips.dualstack.us-west-1.amazonaws.com<br>s3-control-fips.us-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | s3-control.us-west-2.amazonaws.com<br>s3-control-fips.us-west-2.amazonaws.com<br>s3-control.dualstack.us-west-2.amazonaws.com<br>s3-control-fips.dualstack.us-west-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Africa (Cape Town) | af-south-1 | s3-control.af-south-1.amazonaws.com<br>s3-control.dualstack.af-south-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | s3-control.ap-east-1.amazonaws.com<br>s3-control.dualstack.ap-east-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | s3-control.ap-south-2.amazonaws.com<br>s3-control.dualstack.ap-south-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | s3-control.ap-southeast-3.amazonaws.com<br>s3-control.dualstack.ap-southeast-3.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | s3-control.ap-southeast-4.amazonaws.com<br>s3-control.dualstack.ap-southeast-4.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | s3-control.ap-south-1.amazonaws.com<br>s3-control.dualstack.ap-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | s3-control.ap-northeast-3.amazonaws.com<br>s3-control.dualstack.ap-northeast-3.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | s3-control.ap-northeast-2.amazonaws.com<br>s3-control.dualstack.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | s3-control.ap-southeast-1.amazonaws.com<br>s3-control.dualstack.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | s3-control.ap-southeast-2.amazonaws.com<br>s3-control.dualstack.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | s3-control.ap-northeast-1.amazonaws.com<br>s3-control.dualstack.ap-northeast-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | s3-control.ca-central-1.amazonaws.com<br>s3-control.dualstack.ca-central-1.amazonaws.com<br>s3-control-fips.ca-central-1.amazonaws.com<br>s3-control-fips.dualstack.ca-central-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | s3-control.ca-west-1.amazonaws.com<br>s3-control-fips.ca-west-1.amazonaws.com<br>s3-control.dualstack.ca-west-1.amazonaws.com<br>s3-control-fips.dualstack.ca-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | s3-control.eu-central-1.amazonaws.com<br>s3-control.dualstack.eu-central-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | s3-control.eu-west-1.amazonaws.com<br>s3-control.dualstack.eu-west-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | s3-control.eu-west-2.amazonaws.com<br>s3-control.dualstack.eu-west-2.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Milan) | eu-south-1 | s3-control.eu-south-1.amazonaws.com<br>s3-control.dualstack.eu-south-1.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Paris) | eu-west-3 | s3-control.eu-west-3.amazonaws.com<br>s3-control.dualstack.eu-west-3.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Spain) | eu-south-2 | s3-control.eu-south-2.amazonaws.com<br>s3-control.dualstack.eu-south-2.amazonaws.com | HTTPS<br>HTTPS | |
| Europe (Stockholm) | eu-north-1 | s3-control.eu-north-1.amazonaws.com<br>s3-control.dualstack.eu-north-1.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | s3-control.eu-central-2.amazonaws.com | HTTPS |
| | | s3-control.dualstack.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | s3-control.il-central-1.amazonaws.com | HTTPS |
| | | s3-control.dualstack.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | s3-control.me-south-1.amazonaws.com | HTTPS |
| | | s3-control.dualstack.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | s3-control.me-central-1.amazonaws.com | HTTPS |
| | | s3-control.dualstack.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | s3-control.sa-east-1.amazonaws.com | HTTPS |
| | | s3-control.dualstack.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | s3-control.us-gov-east-1.amazonaws.com | HTTPS |
| | | s3-control-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | s3-control.dualstack.us-gov-east-1.amazonaws.com | HTTPS |
| | | s3-control-fips.dualstack.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | s3-control.us-gov-west-1.amazonaws.com | HTTPS |
| | | s3-control-fips.dualstack.us-gov-west-1.amazonaws.com | HTTPS |
| | | s3-control.dualstack.us-gov-west-1.amazonaws.com | HTTPS |
| | | s3-control-fips.us-gov-west-1.amazonaws.com | HTTPS |

**Note**

Amazon S3 dual-stack endpoints support requests to S3 buckets over IPv6 and IPv4. For more information, see Using Dual-Stack Endpoints.

When using the preceding endpoints, the following additional considerations apply:

- Amazon S3 renamed the US Standard Region to the US East (N. Virginia) Region to be consistent with AWS Regional naming conventions. There is no change to the endpoint, and you do not need to make any changes to your application.
- If you use a Region other than the US East (N. Virginia) endpoint to create a bucket, you must set the `LocationConstraint` bucket parameter to the same Region. Both the AWS SDK for Java and AWS SDK for .NET use an enumeration for setting location constraints (`Region` for Java, `S3Region` for .NET). For more information, see `Create Bucket` in the *Amazon Simple Storage Service API Reference*.

### Multivalue answer DNS routing

Amazon S3 uses Multivalue answer (MVA) DNS routing to return multiple answers in response to DNS queries for its endpoints. This behavior improves performance and availability by enabling applications to parallelize their requests across multiple Amazon S3 IP addresses.

To automatically benefit from MVA, you can use the newer versions of the AWS SDKs including the SDK for Java 2.x, SDK for C++, and SDK for Python (Boto3). If you are using the SDK for Java 2.x

with the Common Runtime (CRT) library enabled for asynchronous operations, your requests will automatically use MVA routing.

The AWS CRT-based S3 client is an alternative Amazon S3 asynchronous client that allows you to transfer objects to and from S3 by automatically leveraging the S3 multipart upload API and byte-range fetches. The AWS CRT S3-based client improves reliability by automatically retrying individual failed parts of a multipart file transfer without restarting the transfer from the beginning. When used with MVA, these retries are made to alternative IP addresses, further improving the reliability of your connection to S3.

In cases where applications don't support MVA routing, they will use the first IP from the response. Amazon S3 MVA routing is supported in all AWS Regions, excluding AWS GovCloud (US). For more information, see Multivalue answer routing.

### Amazon S3 website endpoints

When you configure your bucket as a website, the website is available by using the following Region-specific website endpoints. Note that the website endpoints are different than the REST API endpoints listed in the preceding table. For more information about hosting websites on Amazon S3, see Hosting Websites on Amazon S3 in the *Amazon Simple Storage Service User Guide*. You need the hosted zone IDs when using the Amazon Route 53 API to add an alias record to your hosted zone.

> **Note**
>
> Amazon S3 website endpoints do not support HTTPS or Amazon S3 Access Points. If you want to use HTTPS, you can use Amazon CloudFront to serve a static website hosted on Amazon S3. For more information, see Configuring a static website using a custom domain registered with Route 53 and Improving the performance of your website using CloudFront in the *Amazon S3 User Guide*.

### Amazon S3 website endpoints and HostedZone IDs

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | s3-website.us-east-2.amazonaws.com | HTTP | Z2O1EMRO9K5GLX |
| US East (N. Virginia) | us-east-1 | s3-website-us-east-1.amazonaws.com | HTTP | Z3AQBSTGFYJSTF |
| US West (N. California) | us-west-1 | s3-website-us-west-1.amazonaws.com | HTTP | Z2F56UZL2M1ACD |
| US West (Oregon) | us-west-2 | s3-website-us-west-2.amazonaws.com | HTTP | Z3BJ6K6RIION7M |
| Africa (Cape Town) | af-south-1 | s3-website.af-south-1.amazonaws.com | HTTP | Z2OSFR5PIJ8TYW |
| Asia Pacific (Hong Kong) | ap-east-1 | s3-website.ap-east-1.amazonaws.com | HTTP | ZNB98KWMFR0R6 |
| Asia Pacific (Hyderabad) | ap-south-2 | s3-website.ap-south-2.amazonaws.com | HTTP | Z02976202B4EZMXIPMXF7 |
| Asia Pacific (Jakarta) | ap-southeast-3 | s3-website.ap-southeast-3.amazonaws.com | HTTP | Z01846753K324LI26A3VV |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | s3-website.ap-southeast-5.amazonaws.com | HTTP | Z08660063OXLMA7F1FJHU |
| Asia Pacific (Melbourne) | ap-southeast-4 | s3-website.ap-southeast-4.amazonaws.com | HTTP | Z0312387243XT5FE14WFO |
| Asia Pacific (Mumbai) | ap-south-1 | s3-website.ap-south-1.amazonaws.com | HTTP | Z11RGJOFQNVJUP |
| Asia Pacific (New Zealand) | ap-southeast-6 | s3-website.ap-southeast-6.amazonaws.com | HTTP | Z05686083R66JX5C163TC |
| Asia Pacific (Osaka) | ap-northeast-3 | s3-website.ap-northeast-3.amazonaws.com | HTTP | Z2YQB5RD63NC85 |
| Asia Pacific (Seoul) | ap-northeast-2 | s3-website.ap-northeast-2.amazonaws.com | HTTP | Z3W03O7B5YMIYP |
| Asia Pacific (Singapore) | ap-southeast-1 | s3-website-ap-southeast-1.amazonaws.com | HTTP | Z3O0J2DXBE1FTB |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | s3-website-ap-southeast-2.amazonaws.com | HTTP | Z1WCIGYICN2BYD |
| Asia Pacific (Taipei) | ap-east-2 | s3-website.ap-east-2.amazonaws.com | HTTP | Z064739330DAH7WJVOO93 |
| Asia Pacific (Thailand) | ap-southeast-7 | s3-website.ap-southeast-7.amazonaws.com | HTTP | Z0031014GXUMRZG6I14G |
| Asia Pacific (Tokyo) | ap-northeast-1 | s3-website-ap-northeast-1.amazonaws.com | HTTP | Z2M4EHUR26P7ZW |
| Canada (Central) | ca-central-1 | s3-website.ca-central-1.amazonaws.com | HTTP | Z1QDHH18159H29 |
| Canada West (Calgary) | ca-west-1 | s3-website.ca-west-1.amazonaws.com | HTTP | Z03565811Z33SLEZTHOUL |
| Europe (Frankfurt) | eu-central-1 | s3-website.eu-central-1.amazonaws.com | HTTP | Z21DNDUVLTQW6Q |
| Europe (Ireland) | eu-west-1 | s3-website-eu-west-1.amazonaws.com | HTTP | Z1BKCTXD74EZPE |
| Europe (London) | eu-west-2 | s3-website.eu-west-2.amazonaws.com | HTTP | Z3GKZC51ZF0DB4 |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | s3-website.eu-south-1.amazonaws.com | HTTP | Z2OPA49AB41N7K |
| Europe (Paris) | eu-west-3 | s3-website.eu-west-3.amazonaws.com | HTTP | Z3R1K369G5AVDG |
| Europe (Spain) | eu-south-2 | s3-website.eu-south-2.amazonaws.com | HTTP | Z0081959F7139GRJC19J |
| Europe (Stockholm) | eu-north-1 | s3-website.eu-north-1.amazonaws.com | HTTP | Z3BAZG2TWCNX0D |
| Europe (Zurich) | eu-central-2 | s3-website.eu-central-2.amazonaws.com | HTTP | Z030506016YDQGETNASS |
| Israel (Tel Aviv) | il-central-1 | s3-website.il-central-1.amazonaws.com | HTTP | Z09640613K4A3MN55U7GU |
| Mexico (Central) | mx-central-1 | s3-website.mx-central-1.amazonaws.com | HTTP | Z057606446ZNVQJJ8WOP |
| Middle East (Bahrain) | me-south-1 | s3-website.me-south-1.amazonaws.com | HTTP | Z1MPMWCPA7YB62 |
| Middle East (UAE) | me-central-1 | s3-website.me-central-1.amazonaws.com | HTTP | Z06143092I8HRXZRUZROF |

| Region Name | Region | Endpoint | Protocol | Route 53 Hosted Zone ID |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | s3-website-sa-east-1.amazonaws.com | HTTP | Z7KQH4QJS55SO |
| AWS GovCloud (US-East) | us-gov-east-1 | s3-website.us-gov-east-1.amazonaws.com | HTTP | Z2NIFVYYW2VKV1 |
| AWS GovCloud (US-West) | us-gov-west-1 | s3-website-us-gov-west-1.amazonaws.com | HTTP | Z31GFT0UA1I2HV |

# Service quotas

## Amazon S3

> **Note**
>
> The **Default** column shows "Each Account" for all quotas, but some quotas apply per bucket. See the **Description** column for the actual scope.

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Access Points | Each Account: 10,000 | Yes | The number of Amazon S3 Access Points that you can create per region in an account |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Bucket policy | Each Account: 20 Kilobytes | No | The maximum size (in KB) of a bucket policy for an Amazon S3 bucket |
| Bucket tags | Each Account: 50 | No | The maximum number of tags you can assign to an Amazon S3 bucket |
| Directory buckets | Each Account: 100 | Yes | The number of Amazon S3 directory buckets that you can create in an account |
| Event notifications | Each Account: 100 | No | The maximum number of event notifications per Amazon S3 bucket |
| General purpose buckets | Each Account: 10,000 | Yes | The number of Amazon S3 general purpose buckets that you can create in an account |
| Lifecycle rules | Each Account: 1,000 | No | The maximum number of rules you can specify for an Amazon S3 lifecycle configuration |
| Maximum part size | Each Account: 5 Gigabytes | No | The maximum size (in GB) of an Amazon S3 object part in a Multipart upload using the API |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Minimum part size | Each Account: 5 Megabytes | No | The minimum size (in MB) of an Amazon S3 object part in a Multipart upload using the API. The last part uploaded can be less than the stated minimum |
| Multi-Region Access Point Regions | Each Account: 20 | No | Maximum Regions per Multi-Region Access Point. |
| Multi-Region Access Points | Each Account: 100 | No | Multi-Region Access Points per Account. |
| Object size | Each Account: 48.828125 Terabytes | No | The maximum size (in TB) of an Amazon S3 object |
| Object size (Console upload) | Each Account: 160 Gigabytes | No | The maximum size (in GB) of an Amazon S3 object that you can upload using the console |
| Object tags | Each Account: 10 | No | The maximum number of tags you can assign to an Amazon S3 object |
| Parts | Each Account: 10,000 | No | The maximum number of Amazon S3 object parts per Multipart upload |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Replication Destinations | Each Account: 28 | Yes | The number of destination S3 buckets you can add for each source S3 bucket across all replication rules. |
| Replication rules | Each Account: 1,000 | No | The maximum number of rules you can specify in an Amazon S3 Replication configuration |
| Replication transfer rate | Each Account: 1 Gigabits per second | Yes | The maximum Replication Time Control transfer rate that you can replicate from the source region in this account. |
| S3 Glacier: Number of random restore requests. | Each Account: 35 | No | The number of random restore requests from S3 Glacier storage class per PiB stored per day. |
| S3 Glacier: Provisioned capacity units | Each Account: 2 | No | The maximum number of S3 Glacier storage class provisioned capacity units available to purchase per account. |
| S3 Glacier: Retrieval request rate per second. | Each Account: 1,000 | No | The maximum number of retrieval requests per second allowed per account. |

### Amazon S3 on Outposts

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Access Points | Each Account: 10 | No | The maximum number of Amazon S3 on Outposts access points that you can create per S3 on Outposts bucket in the AWS account in the current Outpost. |
| Buckets | Each Account: 100 | No | The maximum number of Amazon S3 on Outposts buckets that you can create per AWS account in the current Outpost. |

## Amazon Glacier endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | glacier.us-east-2.amazonaws.com<br>glacier-fips.us-east-2.amazonaws.com | HTTP and HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | glacier.us-east-1.amazonaws.com<br>glacier-fips.us-east-1.amazonaws.com | HTTP and HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | glacier.us-west-1.amazonaws.com<br>glacier-fips.us-west-1.amazonaws.com | HTTP and HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | glacier.us-west-2.amazonaws.com<br>glacier-fips.us-west-2.amazonaws.com | HTTP and HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | glacier.af-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | glacier.ap-east-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | glacier.ap-southeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | glacier.ap-south-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | glacier.ap-northeast-3.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | glacier.ap-northeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | glacier.ap-southeast-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | glacier.ap-southeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | glacier.ap-northeast-1.amazonaws.com | HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | glacier.ca-central-1.amazonaws.com<br>glacier-fips.ca-central-1.amazonaws.com | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | glacier.eu-central-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | glacier.eu-west-1.amazonaws.com | HTTP and HTTPS | |
| Europe (London) | eu-west-2 | glacier.eu-west-2.amazonaws.com | HTTP and HTTPS | |
| Europe (Milan) | eu-south-1 | glacier.eu-south-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | glacier.eu-west-3.amazonaws.com | HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | glacier.eu-north-1.amazonaws.com | HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | glacier.me-south-1.amazonaws.com | HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | glacier.sa-east-1.amazonaws.com | HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | glacier.us-gov-east-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | glacier.us-gov-west-1.amazonaws.com | HTTP and HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Archive size in GB. | Each supported Region: 40,000 Gigabytes | No | The maximum size of an archive. |
| Archive size. | Each supported Region: 4 Megabytes | No | The minimum size (in MB) of an archive (or part). |
| Multipart parts size. | Each supported Region: 4 Gigabytes | No | The maximum size (in GB) of parts allowed in a multipart upload. |
| Number of multipart parts. | Each supported Region: 10,000 | No | The maximum number of parts allowed in a multipart upload. |
| Number of random restore requests. | Each supported Region: 35 | No | The number of random restore requests per PiB stored per day. |
| Number of vault tags. | Each supported Region: 50 | No | The maximum number of tags that can be applied to a vault. |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Provisioned capacity units | Each supported Region: 2 | No | The maximum number of provisioned capacity units available to purchase per account. |
| Vaults per account | Each supported Region: 1,000 | No | The maximum number of vaults an account can have. |

# Amazon Simple Email Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service API endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | email.us-east-2.amazonaws.com | HTTPS | |
| | | email-fips.us-east-2.api.aws | HTTPS | |
| | | email-fips.us-east-2.amazonaws.com | HTTPS | |
| | | email.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | email.us-east-1.amazonaws.com | HTTPS |
| | | email-fips.us-east-1.api.aws | HTTPS |
| | | email-fips.us-east-1.amazonaws.com | HTTPS |
| | | email.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | email.us-west-1.amazonaws.com | HTTPS |
| | | email-fips.us-west-1.api.aws | HTTPS |
| | | email-fips.us-west-1.amazonaws.com | HTTPS |
| | | email.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | email.us-west-2.amazonaws.com | HTTPS |
| | | email-fips.us-west-2.api.aws | HTTPS |
| | | email-fips.us-west-2.amazonaws.com | HTTPS |
| | | email.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | email.af-south-1.amazonaws.com | HTTPS |
| | | email.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | email.ap-south-2.amazonaws.com | HTTPS |
| | | email.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | email.ap-southeast-3.amazonaws.com | HTTPS |
| | | email.ap-southeast-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | email.ap-southeast-5.amazonaws.com<br>email.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | email.ap-south-1.amazonaws.com<br>email.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | email.ap-northeast-3.amazonaws.com<br>email.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | email.ap-northeast-2.amazonaws.com<br>email.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | email.ap-southeast-1.amazonaws.com<br>email.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | email.ap-southeast-2.amazonaws.com<br>email.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | email.ap-northeast-1.amazonaws.com<br>email.ap-northeast-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | email.ca-central-1.amazonaws.com | HTTPS | |
| | | email-fips.ca-central-1.api.aws | HTTPS | |
| | | email-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | email.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | email.ca-west-1.amazonaws.com | HTTPS | |
| | | email-fips.ca-west-1.api.aws | HTTPS | |
| | | email-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | email.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | email.eu-central-1.amazonaws.com | HTTPS | |
| | | email.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | email.eu-west-1.amazonaws.com | HTTPS | |
| | | email.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | email.eu-west-2.amazonaws.com | HTTPS | |
| | | email.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | email.eu-south-1.amazonaws.com | HTTPS | |
| | | email.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | email.eu-west-3.amazonaws.com | HTTPS | |
| | | email.eu-west-3.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | email.eu-north-1.amazonaws.com | HTTPS | |
| | | email.eu-north-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | email.eu-central-2.amazonaws.com | HTTPS |
| | | email.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | email.il-central-1.amazonaws.com | HTTPS |
| | | email.il-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | email.me-south-1.amazonaws.com | HTTPS |
| | | email.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | email.me-central-1.amazonaws.com | HTTPS |
| | | email.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | email.sa-east-1.amazonaws.com | HTTPS |
| | | email.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | email.us-gov-east-1.amazonaws.com | HTTPS |
| | | email-fips.us-gov-east-1.api.aws | HTTPS |
| | | email-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | email.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | email.us-gov-west-1.amazonaws.com | HTTPS |
| | | email-fips.us-gov-west-1.api.aws | HTTPS |
| | | email-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | email.us-gov-west-1.api.aws | HTTPS |

# SMTP endpoints

> **Note**
>
> SMTP endpoints are not currently available in Africa (Cape Town), Asia Pacific (Hyderabad), Asia Pacific (Jakarta), Europe (Milan), Europe (Zurich), Israel (Tel Aviv), Middle East (Bahrain), Middle East (UAE), Canada West (Calgary), Asia Pacific (Malaysia).

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | email-smtp.us-east-2.amazonaws.com<br>email-smtp.us-east-2.api.aws<br>email-smtp-fips.us-east-2.amazonaws.com<br>email-smtp-fips.us-east-2.api.aws | SMTP |
| US East (N. Virginia) | us-east-1 | email-smtp.us-east-1.amazonaws.com<br>email-smtp.us-east-1.api.aws<br>email-smtp-fips.us-east-1.amazonaws.com<br>email-smtp-fips.us-east-1.api.aws | SMTP |
| US West (N. California) | us-west-1 | email-smtp.us-west-1.amazonaws.com | SMTP |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | email-smtp.us-west-1.api.aws<br><br>email-smtp-fips.us-west-1.amazonaws.com<br><br>email-smtp-fips.us-west-1.api.aws | |
| US West (Oregon) | us-west-2 | email-smtp.us-west-2.amazonaws.com<br><br>email-smtp.us-west-2.api.aws<br><br>email-smtp-fips.us-west-2.amazonaws.com<br><br>email-smtp-fips.us-west-2.api.aws | SMTP |
| Asia Pacific (Mumbai) | ap-south-1 | email-smtp.ap-south-1.amazonaws.com<br><br>email-smtp.ap-south-1.api.aws | SMTP |
| Asia Pacific (Osaka) | ap-northeast-3 | email-smtp.ap-northeast-3.amazonaws.com<br><br>email-smtp.ap-northeast-3.api.aws | SMTP |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | email-smtp.ap-northeast-2.amazonaws.com<br><br>email-smtp.ap-northeast-2.api.aws | SMTP |
| Asia Pacific (Singapore) | ap-southeast-1 | email-smtp.ap-southeast-1.amazonaws.com<br><br>email-smtp.ap-southeast-1.api.aws | SMTP |
| Asia Pacific (Sydney) | ap-southeast-2 | email-smtp.ap-southeast-2.amazonaws.com<br><br>email-smtp.ap-southeast-2.api.aws | SMTP |
| Asia Pacific (Tokyo) | ap-northeast-1 | email-smtp.ap-northeast-1.amazonaws.com<br><br>email-smtp.ap-northeast-1.api.aws | SMTP |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | email-smtp.ca-central-1.amazonaws.com<br>email-smtp.ca-central-1.api.aws<br>email-smtp-fips.ca-central-1.amazonaws.com<br>email-smtp-fips.ca-central-1.api.aws | SMTP |
| Europe (Frankfurt) | eu-central-1 | email-smtp.eu-central-1.amazonaws.com<br>email-smtp.eu-central-1.api.aws | SMTP |
| Europe (Ireland) | eu-west-1 | email-smtp.eu-west-1.amazonaws.com<br>email-smtp.eu-west-1.api.aws | SMTP |
| Europe (London) | eu-west-2 | email-smtp.eu-west-2.amazonaws.com<br>email-smtp.eu-west-2.api.aws | SMTP |
| Europe (Paris) | eu-west-3 | email-smtp.eu-west-3.amazonaws.com<br>email-smtp.eu-west-3.api.aws | SMTP |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| Europe (Stockholm) | eu-north-1 | email-smtp.eu-north-1.amazonaws.com<br><br>email-smtp.eu-north-1.api.aws | SMTP |
| South America (São Paulo) | sa-east-1 | email-smtp.sa-east-1.amazonaws.com<br><br>email-smtp.sa-east-1.api.aws | SMTP |
| AWS GovCloud (US-West) | us-gov-west-1 | email-smtp.us-gov-west-1.amazonaws.com<br><br>email-smtp.us-gov-west-1.api.aws<br><br>email-smtp-fips.us-gov-west-1.amazonaws.com<br><br>email-smtp-fips.us-gov-west-1.api.aws | SMTP |

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| AWS GovCloud (US-East) | us-gov-east-1 | email-smtp.us-gov-east-1.amazonaws.com<br><br>email-smtp.us-gov-east-1.api.aws<br><br>email-smtp-fips.us-gov-east-1.amazonaws.com<br><br>email-smtp-fips.us-gov-east-1.api.aws | SMTP |

## Email Receiving endpoints

| Region Name | Region | Email Receiving Endpoint |
| --- | --- | --- |
| US East (N. Virginia) | us-east-1 | inbound-smtp.us-east-1.amazonaws.com |
| US East (Ohio) | us-east-2 | inbound-smtp.us-east-2.amazonaws.com |
| US West (N. California) | us-west-1 | inbound-smtp.us-west-1.amazonaws.com |
| US West (Oregon) | us-west-2 | inbound-smtp.us-west-2.amazonaws.com |
| Africa (Cape Town) | af-south-1 | inbound-smtp.af-south-1.amazonaws.com |
| Asia Pacific (Jakarta) | ap-southeast-3 | inbound-smtp.ap-southeast-3.amazonaws.com |

| Region Name | Region | Email Receiving Endpoint |
|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | inbound-smtp.ap-south-1.amazonaws.com |
| Asia Pacific (Osaka) | ap-northeast-3 | inbound-smtp.ap-northeast-3.amazonaws.com |
| Asia Pacific (Seoul) | ap-northeast-2 | inbound-smtp.ap-northeast-2.amazonaws.com |
| Asia Pacific (Singapore) | ap-southeast-1 | inbound-smtp.ap-southeast-1.amazonaws.com |
| Asia Pacific (Sydney) | ap-southeast-2 | inbound-smtp.ap-southeast-2.amazonaws.com |
| Asia Pacific (Tokyo) | ap-northeast-1 | inbound-smtp.ap-northeast-1.amazonaws.com |
| Canada (Central) | ca-central-1 | inbound-smtp.ca-central-1.amazonaws.com |
| Europe (Frankfurt) | eu-central-1 | inbound-smtp.eu-central-1.amazonaws.com |
| Europe (Ireland) | eu-west-1 | inbound-smtp.eu-west-1.amazonaws.com |
| Europe (London) | eu-west-2 | inbound-smtp.eu-west-2.amazonaws.com |
| Europe (Milan) | eu-south-1 | inbound-smtp.eu-south-1.amazonaws.com |
| Europe (Paris) | eu-west-3 | inbound-smtp.eu-west-3.amazonaws.com |
| Europe (Stockholm) | eu-north-1 | inbound-smtp.eu-north-1.amazonaws.com |

| Region Name | Region | Email Receiving Endpoint |
|---|---|---|
| Israel (Tel Aviv) | il-central-1 | inbound-smtp.il-central-1.amazonaws.com |
| Middle East (Bahrain) | me-south-1 | inbound-smtp.me-south-1.amazonaws.com |
| South America (São Paulo) | sa-east-1 | inbound-smtp.sa-east-1.amazonaws.com |

Amazon SES does not support email receiving in the following Regions: *AWS GovCloud (US-West) and AWS GovCloud (US-East).*.

## DKIM domains

| Region Name | Region | SES DKIM Domain |
|---|---|---|
| Africa (Cape Town) | af-south-1 | dkim.af-south-1.amazonses.com |
| Asia Pacific (Hyderabad) | ap-south-2 | dkim.ap-south-2.amazonses.com |
| Asia Pacific (Jakarta) | ap-southeast-3 | dkim.ap-southeast-3.amazonses.com |
| Asia Pacific (Malaysia) | ap-southeast-5 | dkim.ap-southeast-5.amazonses.com |
| Canada West (Calgary) | ca-west-1 | dkim.ca-west-1.amazonses.com |
| Asia Pacific (Osaka) | ap-northeast-3 | dkim.ap-northeast-3.amazonses.com |
| Europe (Milan) | eu-south-1 | dkim.eu-south-1.amazonses.com |

| Region Name | Region | SES DKIM Domain |
|---|---|---|
| Europe (Zurich) | eu-central-2 | dkim.eu-central-2.amazonses.com |
| Israel (Tel Aviv) | il-central-1 | dkim.il-central-1.amazonses.com |
| Middle East (UAE) | me-central-1 | dkim.me-central-1.amazonses.com |
| AWS GovCloud (US-East) | us-gov-east-1 | dkim.us-gov-east-1.amazonses.com |
| All other regions | | dkim.amazonses.com |

## Tracking domains for open/click links (HTTP & HTTPS)

| Region Name | Tracking Domain |
|---|---|
| US East (Ohio) | `r.us-east-2.awstrack.me` |
| US East (N. Virginia) | `r.us-east-1.awstrack.me` |
| US West (N. California) | `r.us-west-1.awstrack.me` |
| US West (Oregon) | `r.us-west-2.awstrack.me` |
| Africa (Cape Town) | `r.af-south-1.awstrack.me` |
| Asia Pacific (Hyderabad) | `r.ap-south-2.awstrack.me` |
| Asia Pacific (Jakarta) | `r.ap-southeast-3.awstrack.me` |
| Asia Pacific (Malaysia) | `r.ap-southeast-5.awstrack.me` |
| Asia Pacific (Mumbai) | `r.ap-south-1.awstrack.me` |
| Asia Pacific (Osaka) | `r.ap-northeast-3.awstrack.me` |

| Region Name | Tracking Domain |
| --- | --- |
| Asia Pacific (Seoul) | r.ap-northeast-2.awstrack.me |
| Asia Pacific (Singapore) | r.ap-southeast-1.awstrack.me |
| Asia Pacific (Sydney) | r.ap-southeast-2.awstrack.me |
| Asia Pacific (Tokyo) | r.ap-northeast-1.awstrack.me |
| Canada (Central) | r.ca-central-1.awstrack.me |
| Canada West (Calgary) | r.ca-west-1.awstrack.me |
| Europe (Frankfurt) | r.eu-central-1.awstrack.me |
| Europe (Ireland) | r.eu-west-1.awstrack.me |
| Europe (London) | r.eu-west-2.awstrack.me |
| Europe (Milan) | r.eu-south-1.awstrack.me |
| Europe (Stockholm) | r.eu-north-1.awstrack.me |
| Europe (Zurich) | r.eu-central-2.awstrack.me |
| Israel (Tel Aviv) | r.il-central-1.awstrack.me |
| Middle East (Bahrain) | r.me-south-1.awstrack.me |
| Middle East (UAE) | r.me-central-1.awstrack.me |
| South America (São Paulo) | r.sa-east-1.awstrack.me |
| AWS GovCloud (US-West) | r.us-gov-west-1.awstrack.me |
| AWS GovCloud (US-East) | r.us-gov-east-1.awstrack.me |

# Feedback endpoints used by SES for Custom MAIL FROM domains

| Region Name | Endpoint |
|---|---|
| US East (Ohio) | feedback-smtp.us-east-2.amazonses.com |
| US East (N. Virginia) | feedback-smtp.us-east-1.amazonses.com |
| US West (N. California) | feedback-smtp.us-west-1.amazonses.com |
| US West (Oregon) | feedback-smtp.us-west-2.amazonses.com |
| Africa (Cape Town) | feedback-smtp.af-south-1.amazonses.com |
| Asia Pacific (Hyderabad) | feedback-smtp.ap-south-2.amazonses.com |
| Asia Pacific (Jakarta) | feedback-smtp.ap-southeast-3.amazonses.com |
| Asia Pacific (Mumbai) | feedback-smtp.ap-south-1.amazonses.com |
| Asia Pacific (Osaka) | feedback-smtp.ap-northeast-3.amazonses.com |
| Asia Pacific (Seoul) | feedback-smtp.ap-northeast-2.amazonses.com |
| Asia Pacific (Singapore) | feedback-smtp.ap-southeast-1.amazonses.com |
| Asia Pacific (Sydney) | feedback-smtp.ap-southeast-2.amazonses.com |
| Asia Pacific (Malaysia) | feedback-smtp.ap-southeast-5.amazonses.com |
| Asia Pacific (Tokyo) | feedback-smtp.ap-northeast-1.amazonses.com |
| Canada (Central) | feedback-smtp.ca-central-1.amazonses.com |

| Region Name | Endpoint |
| --- | --- |
| Canada West (Calgary) | feedback-smtp.ca-west-1.amazonses.com |
| Europe (Frankfurt) | feedback-smtp.eu-central-1.amazonses.com |
| Europe (Ireland) | feedback-smtp.eu-west-1.amazonses.com |
| Europe (London) | feedback-smtp.eu-west-2.amazonses.com |
| Europe (Milan) | feedback-smtp.eu-south-1.amazonses.com |
| Europe (Paris) | feedback-smtp.eu-west-3.amazonses.com |
| Europe (Stockholm) | feedback-smtp.eu-north-1.amazonses.com |
| Europe (Zurich) | feedback-smtp.eu-central-2.amazonses.com |
| Israel (Tel Aviv) | feedback-smtp.il-central-1.amazonses.com |
| Middle East (Bahrain) | feedback-smtp.me-south-1.amazonses.com |
| Middle East (UAE) | feedback-smtp.me-central-1.amazonses.com |
| South America (São Paulo) | feedback-smtp.sa-east-1.amazonses.com |
| AWS GovCloud (US-West) | feedback-smtp.us-gov-west-1.amazonses.com |
| AWS GovCloud (US-East) | feedback-smtp.us-gov-east-1.amazonses.com |

## IP ranges used by Mail Manager SMTP relays

| Region | IP range |
| --- | --- |
| us-east-1 | 206.55.129.0/24 |
| us-east-2 | 206.55.134.0/24 |
| us-west-1 | 206.55.142.0/24 |

| Region | IP range |
| --- | --- |
| us-west-2 | 206.55.131.0/24 |
| af-south-1 | 216.221.183.0/24 |
| ap-northeast-1 | 206.55.130.0/24 |
| ap-northeast-2 | 206.55.139.0/24 |
| ap-northeast-3 | 206.55.140.0/24 |
| ap-south-1 | 206.55.136.0/24 |
| ap-south-2 | 216.221.185.0/24 |
| ap-southeast-1 | 206.55.143.0/24 |
| ap-southeast-2 | 206.55.128.0/24 |
| ap-southeast-3 | 216.221.180.0/24 |
| ap-southeast-5 | 216.221.178.0/24 |
| ca-central-1 | 216.221.176.0/24 |
| ca-west-1 | 216.221.186.0/24 |
| il-central-1 | 216.221.181.0/24 |
| eu-central-1 | 206.55.132.0/24 |
| eu-central-2 | 216.221.184.0/24 |
| eu-north-1 | 206.55.135.0/24 |
| eu-south-1 | 216.221.182.0/24 |
| eu-west-1 | 206.55.133.0/24 |
| eu-west-2 | 206.55.141.0/24 |

| Region | IP range |
| --- | --- |
| eu-west-3 | 206.55.137.0/24 |
| me-central-1 | 216.221.177.0/24 |
| me-south-1 | 216.221.179.0/24 |
| sa-east-1 | 206.55.138.0/24 |
| us-gov-east-1 | 216.221.188.0/24 |
| us-gov-west-1 | 216.221.187.0/24 |

## Service quotas

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Configuration set count | Each supported Region: 10,000 | Yes | The maximum number of configuration sets that can exist within the current Region. |
| Sending quota | Each supported Region: 200 per 24 hours | Yes | The maximum number of emails that you can send in a 24-hour period for this account in the current Region. |
| Sending rate | Each supported Region: 1 per second | Yes | The maximum number of emails that Amazon SES can accept each second for this account in the current Region. |
| Tenant count | Each supported Region: 10,000 | Yes | The maximum number of tenants that can |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| | | | exist within the current Region. |

For more information, see Service quotas in Amazon SES in the *Amazon Simple Email Service Developer Guide*.

# AWS Signer endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints supporting signing for AWS Lambda, Amazon ECR, and Amazon EKS

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | signer.us-east-2.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-east-2.api.aws | HTTPS | |
| | | data-signer.us-east-2.api.aws | HTTPS | |
| | | signer-fips.us-east-2.api.aws | HTTPS | |
| | | signer-fips.us-east-2.amazonaws.com | HTTPS | |
| | | signer.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | data-signer.us-east-2.amazonaws.com | HTTPS | |
| | | verification.signer-fips.us-east-2.amazonaws.com | HTTPS | |
| | | verification.signer.us-east-2.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | signer.us-east-1.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-east-1.api.aws | HTTPS | |
| | | data-signer.us-east-1.api.aws | HTTPS | |
| | | signer-fips.us-east-1.api.aws | HTTPS | |
| | | signer-fips.us-east-1.amazonaws.com | HTTPS | |
| | | signer.us-east-1.api.aws | HTTPS | |
| | | data-signer.us-east-1.amazonaws.com | HTTPS | |
| | | verification.signer-fips.us-east-1.amazonaws.com | HTTPS | |
| | | verification.signer.us-east-1.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US West (N. California) | us-west-1 | signer.us-west-1.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-west-1.api.aws | HTTPS | |
| | | data-signer.us-west-1.api.aws | HTTPS | |
| | | signer-fips.us-west-1.api.aws | HTTPS | |
| | | signer-fips.us-west-1.amazonaws.com | HTTPS | |
| | | signer.us-west-1.api.aws | HTTPS | |
| | | data-signer.us-west-1.amazonaws.com | HTTPS | |
| | | verification.signer-fips.us-west-1.amazonaws.com | HTTPS | |
| | | verification.signer.us-west-1.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | signer.us-west-2.amazonaws.com | HTTPS |
| | | data-signer-fips.us-west-2.api.aws | HTTPS |
| | | data-signer.us-west-2.api.aws | HTTPS |
| | | signer-fips.us-west-2.api.aws | HTTPS |
| | | signer-fips.us-west-2.amazonaws.com | HTTPS |
| | | signer.us-west-2.api.aws | HTTPS |
| | | data-signer.us-west-2.amazonaws.com | HTTPS |
| | | verification.signer-fips.us-west-2.amazonaws.com | HTTPS |
| | | verification.signer.us-west-2.amazonaws.com | HTTPS |
| | | data-signer-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | signer.af-south-1.amazonaws.com | HTTPS |
| | | data-signer.af-south-1.amazonaws.com | HTTPS |
| | | data-signer.af-south-1.api.aws | HTTPS |
| | | signer.af-south-1.api.aws | HTTPS |
| | | verification.signer.af-south-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | signer.ap-east-1.amazonaws.com | HTTPS |
| | | data-signer.ap-east-1.amazonaws.com | HTTPS |
| | | data-signer.ap-east-1.api.aws | HTTPS |
| | | signer.ap-east-1.api.aws | HTTPS |
| | | verification.signer.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | signer.ap-south-1.amazonaws.com | HTTPS |
| | | data-signer.ap-south-1.amazonaws.com | HTTPS |
| | | data-signer.ap-south-1.api.aws | HTTPS |
| | | signer.ap-south-1.api.aws | HTTPS |
| | | verification.signer.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | signer.ap-northeast-2.amazonaws.com | HTTPS |
| | | data-signer.ap-northeast-2.amazonaws.com | HTTPS |
| | | data-signer.ap-northeast-2.api.aws | HTTPS |
| | | signer.ap-northeast-2.api.aws | HTTPS |
| | | verification.signer.ap-northeast-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | signer.ap-southeast-1.amazonaws.com | HTTPS | |
| | | data-signer.ap-southeast-1.amazonaws.com | HTTPS | |
| | | data-signer.ap-southeast-1.api.aws | HTTPS | |
| | | signer.ap-southeast-1.api.aws | HTTPS | |
| | | verification.signer.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | signer.ap-southeast-2.amazonaws.com | HTTPS | |
| | | data-signer.ap-southeast-2.amazonaws.com | HTTPS | |
| | | data-signer.ap-southeast-2.api.aws | HTTPS | |
| | | signer.ap-southeast-2.api.aws | HTTPS | |
| | | verification.signer.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | signer.ap-northeast-1.amazonaws.com | HTTPS | |
| | | data-signer.ap-northeast-1.amazonaws.com | HTTPS | |
| | | data-signer.ap-northeast-1.api.aws | HTTPS | |
| | | signer.ap-northeast-1.api.aws | HTTPS | |
| | | verification.signer.ap-northeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | signer.ca-central-1.amazonaws.com | HTTPS |
| | | data-signer.ca-central-1.amazonaws.com | HTTPS |
| | | data-signer.ca-central-1.api.aws | HTTPS |
| | | signer.ca-central-1.api.aws | HTTPS |
| | | verification.signer.ca-central-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | signer.eu-central-1.amazonaws.com | HTTPS |
| | | data-signer.eu-central-1.amazonaws.com | HTTPS |
| | | data-signer.eu-central-1.api.aws | HTTPS |
| | | signer.eu-central-1.api.aws | HTTPS |
| | | verification.signer.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | signer.eu-west-1.amazonaws.com | HTTPS |
| | | data-signer.eu-west-1.amazonaws.com | HTTPS |
| | | data-signer.eu-west-1.api.aws | HTTPS |
| | | signer.eu-west-1.api.aws | HTTPS |
| | | verification.signer.eu-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | signer.eu-west-2.amazonaws.com | HTTPS |
| | | data-signer.eu-west-2.amazonaws.com | HTTPS |
| | | data-signer.eu-west-2.api.aws | HTTPS |
| | | signer.eu-west-2.api.aws | HTTPS |
| | | verification.signer.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | signer.eu-south-1.amazonaws.com | HTTPS |
| | | data-signer.eu-south-1.amazonaws.com | HTTPS |
| | | data-signer.eu-south-1.api.aws | HTTPS |
| | | signer.eu-south-1.api.aws | HTTPS |
| | | verification.signer.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | signer.eu-west-3.amazonaws.com | HTTPS |
| | | data-signer.eu-west-3.amazonaws.com | HTTPS |
| | | data-signer.eu-west-3.api.aws | HTTPS |
| | | signer.eu-west-3.api.aws | HTTPS |
| | | verification.signer.eu-west-3.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | signer.eu-north-1.amazonaws.com | HTTPS |
| | | data-signer.eu-north-1.amazonaws.com | HTTPS |
| | | data-signer.eu-north-1.api.aws | HTTPS |
| | | signer.eu-north-1.api.aws | HTTPS |
| | | verification.signer.eu-north-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | signer.me-south-1.amazonaws.com | HTTPS | |
| | | data-signer.me-south-1.amazonaws.com | HTTPS | |
| | | data-signer.me-south-1.api.aws | HTTPS | |
| | | signer.me-south-1.api.aws | HTTPS | |
| | | verification.signer.me-south-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | signer.sa-east-1.amazonaws.com | HTTPS | |
| | | data-signer.sa-east-1.amazonaws.com | HTTPS | |
| | | data-signer.sa-east-1.api.aws | HTTPS | |
| | | signer.sa-east-1.api.aws | HTTPS | |
| | | verification.signer.sa-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | signer.us-gov-east-1.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-gov-east-1.api.aws | HTTPS | |
| | | data-signer.us-gov-east-1.api.aws | HTTPS | |
| | | signer-fips.us-gov-east-1.api.aws | HTTPS | |
| | | signer-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | signer.us-gov-east-1.api.aws | HTTPS | |
| | | data-signer.us-gov-east-1.amazonaws.com | HTTPS | |
| | | verification.signer-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | verification.signer.us-gov-east-1.amazonaws.com | HTTPS | |
| | | data-signer-fips.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | signer.us-gov-west-1.amazonaws.com | HTTPS |
| | | data-signer-fips.us-gov-west-1.api.aws | HTTPS |
| | | data-signer.us-gov-west-1.api.aws | HTTPS |
| | | signer-fips.us-gov-west-1.api.aws | HTTPS |
| | | signer-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | signer.us-gov-west-1.api.aws | HTTPS |
| | | data-signer.us-gov-west-1.amazonaws.com | HTTPS |
| | | verification.signer-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | verification.signer.us-gov-west-1.amazonaws.com | HTTPS |
| | | data-signer-fips.us-gov-west-1.amazonaws.com | HTTPS |

## Service endpoints supporting signing for FreeRTOS and AWS IoT Device Management

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | signer.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | signer.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | signer.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US West (Oregon) | us-west-2 | signer.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | signer.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | signer.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | signer.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | signer.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | signer.ap-southeast-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | signer.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | signer.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | signer.ca-central-1.amazonaws.com | HTTPS |
| China (Beijing) | cn-north-1 | acm.cn-north-1.amazonaws.com.cn | HTTPS |
| China (Ningxia) | cn-northwest-1 | acm.cn-northwest-1.amazonaws.com.cn | HTTPS |
| Europe (Frankfurt) | eu-central-1 | signer.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | signer.eu-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | signer.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | signer.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | signer.eu-west-3.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | signer.eu-north-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | signer.me-south-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | signer.sa-east-1.amazonaws.com | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of AddProfilePermission requests | Each supported Region: 3 | Yes | The maximum number of AddProfilePermission requests that you can make, per second, in this account in the current region. |
| Rate of CancelSigningProfile requests | Each supported Region: 3 | Yes | The maximum number of CancelSigningProfile requests that you can make, per second, in this |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | account in the current region. |
| Rate of DescribeSigningJob requests | Each supported Region: 6 | Yes | The maximum number of DescribeSigningJob requests that you can make, per second, in this account in the current region. |
| Rate of GetRevocationStatus requests | Each supported Region: 100 | Yes | The maximum number of GetRevocationStatus requests that you can make, per second, in this account in the current region. |
| Rate of GetSigningPlatform requests | Each supported Region: 3 | Yes | The maximum number of GetSigningPlatform requests that you can make, per second, in this account in the current region. |
| Rate of GetSigningProfile requests | Each supported Region: 3 | Yes | The maximum number of GetSigningProfile requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of ListProfilePermissions requests | Each supported Region: 6 | Yes | The maximum number of ListProfilePermissions requests that you can make, per second, in this account in the current region. |
| Rate of ListSigningJobs requests | Each supported Region: 6 | Yes | The maximum number of ListSigningJobs requests that you can make, per second, in this account in the current region. |
| Rate of ListSigningPlatforms requests | Each supported Region: 6 | Yes | The maximum number of ListSigningPlatforms requests that you can make, per second, in this account in the current region. |
| Rate of ListSigningProfiles requests | Each supported Region: 6 | Yes | The maximum number of ListSigningProfiles requests that you can make, per second, in this account in the current region. |
| Rate of ListTagsForResource requests | Each supported Region: 6 | Yes | The maximum number of ListTagsForResource requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of PutSigningProfile requests | Each supported Region: 3 | Yes | The maximum number of PutSigningProfile requests that you can make, per second, in this account in the current region. |
| Rate of RemoveProfilePermission requests | Each supported Region: 3 | Yes | The maximum number of RemoveProfilePermission requests that you can make, per second, in this account in the current region. |
| Rate of RevokeSignature requests | Each supported Region: 3 | Yes | The maximum number of RevokeSignature requests that you can make, per second, in this account in the current region. |
| Rate of RevokeSigningProfile requests | Each supported Region: 3 | Yes | The maximum number of RevokeSigningProfile requests that you can make, per second, in this account in the current region. |
| Rate of SignPayload requests | Each supported Region: 10 | Yes | The maximum number of SignPayload requests that you can make, per second, in this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rate of StartSigningJob requests | Each supported Region: 3 | Yes | The maximum number of StartSigningJob requests that you can make, per second, in this account in the current region. |
| Rate of TagResource requests | Each supported Region: 3 | Yes | The maximum number of TagResource requests that you can make, per second, in this account in the current region. |
| Rate of UntagResource requests | Each supported Region: 3 | Yes | The maximum number of UntagResource requests that you can make, per second, in this account in the current region. |

## AWS Sign-In endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | us-east-2.signin.aws.amazon.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | signin.aws.amazon.com | HTTPS | |
| US West (N. California) | us-west-1 | us-west-1.signin.aws.amazon.com | HTTPS | |
| US West (Oregon) | us-west-2 | us-west-2.signin.aws.amazon.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | af-south-1.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | ap-east-1.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | ap-south-2.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ap-southeast-3.signin.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | ap-southeast-5.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ap-southeast-4.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | ap-south-1.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ap-southeast-6.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ap-northeast-3.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | ap-northeast-2.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | ap-southeast-1.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | ap-southeast-2.signin.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | ap-east-2.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | ap-southeast-7.signin.aws.amazon.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | ap-northeast-1.signin.aws.amazon.com | HTTPS | |
| Canada (Central) | ca-central-1 | ca-central-1.signin.aws.amazon.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | ca-west-1.signin.aws.amazon.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | eu-central-1.signin.aws.amazon.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | eu-west-1.signin.aws.amazon.com | HTTPS | |
| Europe (London) | eu-west-2 | eu-west-2.signin.aws.amazon.com | HTTPS | |
| Europe (Milan) | eu-south-1 | eu-south-1.signin.aws.amazon.com | HTTPS | |
| Europe (Paris) | eu-west-3 | eu-west-3.signin.aws.amazon.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Spain) | eu-south-2 | eu-south-2.signin.aws.amazon.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | eu-north-1.signin.aws.amazon.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | eu-central-2.signin.aws.amazon.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | il-central-1.signin.aws.amazon.com | HTTPS | |
| Mexico (Central) | mx-central-1 | mx-central-1.signin.aws.amazon.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | me-south-1.signin.aws.amazon.com | HTTPS | |
| Middle East (UAE) | me-central-1 | me-central-1.signin.aws.amazon.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | sa-east-1.signin.aws.amazon.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | us-gov-east-1.signin.amazonaws-us-gov.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | signin.amazonaws-us-gov.com | HTTPS | |

## OAuth service endpoints

The following table lists the endpoints that you can use for AWS Sign-In OAuth operations.

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | https://us-east-1.oauth.signin.aws | HTTPS |
| US East (Ohio) | us-east-2 | https://us-east-2.oauth.signin.aws | HTTPS |
| US West (N. California) | us-west-1 | https://us-west-1.oauth.signin.aws | HTTPS |
| US West (Oregon) | us-west-2 | https://us-west-2.oauth.signin.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | https://af-south-1.oauth.signin.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | https://ap-east-1.oauth.signin.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | https://ap-south-2.oauth.signin.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | https://ap-southeast-3.oauth.signin.aws | HTTPS |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | https://ap-southeast-4.oauth.signin.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | https://ap-south-1.oauth.signin.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | https://ap-northeast-3.oauth.signin.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | https://ap-northeast-2.oauth.signin.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | https://ap-southeast-1.oauth.signin.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | https://ap-southeast-2.oauth.signin.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | https://ap-northeast-1.oauth.signin.aws | HTTPS |
| Canada (Central) | ca-central-1 | https://ca-central-1.oauth.signin.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | https://ca-west-1.oauth.signin.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | https://eu-central-1.oauth.signin.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | https://eu-west-1.oauth.signin.aws | HTTPS |
| Europe (London) | eu-west-2 | https://eu-west-2.oauth.signin.aws | HTTPS |
| Europe (Milan) | eu-south-1 | https://eu-south-1.oauth.signin.aws | HTTPS |

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | https://eu-west-3.oauth.signin.aws | HTTPS |
| Europe (Spain) | eu-south-2 | https://eu-south-2.oauth.signin.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | https://eu-north-1.oauth.signin.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | https://eu-central-2.oauth.signin.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | https://il-central-1.oauth.signin.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | https://me-south-1.oauth.signin.aws | HTTPS |
| Middle East (UAE) | me-central-1 | https://me-central-1.oauth.signin.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | https://sa-east-1.oauth.signin.aws | HTTPS |

## Service quotas

AWS Sign-In has no increasable quotas.

# Amazon SimpleDB endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | sdb.amazonaws.com | HTTP and HTTPS | |
| US West (N. California) | us-west-1 | sdb.us-west-1.amazonaws.com | HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | sdb.us-west-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | sdb.ap-southeast-1.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | sdb.ap-southeast-2.amazonaws.com | HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | sdb.ap-northeast-1.amazonaws.com | HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | sdb.eu-west-1.amazonaws.com | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | sdb.sa-east-1.amazonaws.com | HTTP and HTTPS | |

### Service quotas

| Resource | Default |
|---|---|
| Domains | 250 |

For more information, see Amazon SimpleDB Quotas in the *Amazon SimpleDB Developer Guide*.

## Amazon Simple Notification Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, IPv6 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

### Service endpoints

Service endpoints allow you to connect to other AWS services. The following endpoints support both standard and FIFO topics.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | sns.us-east-2.amazonaws.com<br>sns-fips.us-east-2.api.aws<br>sns.us-east-2.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US East (N. Virginia) | us-east-1 | sns.us-east-1.amazonaws.com<br>sns-fips.us-east-1.api.aws<br>sns.us-east-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (N. California) | us-west-1 | sns.us-west-1.amazonaws.com<br>sns-fips.us-west-1.api.aws<br>sns.us-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | sns.us-west-2.amazonaws.com<br>sns-fips.us-west-2.api.aws<br>sns.us-west-2.api.aws | HTTP and HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | | HTTP and HTTPS | |
| Africa (Cape Town) | af-south-1 | sns.af-south-1.amazonaws.com<br><br>sns.af-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | sns.ap-east-1.amazonaws.com<br><br>sns.ap-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | sns.ap-south-2.amazonaws.com<br><br>sns.ap-south-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | sns.ap-southeast-3.amazonaws.com<br><br>sns.ap-southeast-3.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | sns.ap-southeast-5.amazonaws.com<br>sns.ap-southeast-5.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | sns.ap-southeast-4.amazonaws.com<br>sns.ap-southeast-4.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | sns.ap-south-1.amazonaws.com<br>sns.ap-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | sns.ap-southeast-6.amazonaws.com<br>sns.ap-southeast-6.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | sns.ap-northeast-3.amazonaws.com<br>sns.ap-northeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | sns.ap-northeast-2.amazonaws.com<br>sns.ap-northeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | sns.ap-southeast-1.amazonaws.com<br>sns.ap-southeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | sns.ap-southeast-2.amazonaws.com<br>sns.ap-southeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | sns.ap-east-2.amazonaws.com<br>sns.ap-east-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | sns.ap-southeast-7.amazonaws.com<br>sns.ap-southeast-7.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | sns.ap-northeast-1.amazonaws.com<br>sns.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | sns.ca-central-1.amazonaws.com<br>sns-fips.ca-central-1.api.aws<br>sns.ca-central-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | sns.ca-west-1.amazonaws.com<br>sns-fips.ca-west-1.api.aws<br>sns.ca-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | sns.eu-central-1.amazonaws.com<br>sns.eu-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | sns.eu-west-1.amazonaws.com<br>sns.eu-west-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (London) | eu-west-2 | sns.eu-west-2.amazonaws.com<br>sns.eu-west-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | sns.eu-south-1.amazonaws.com<br><br>sns.eu-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | sns.eu-west-3.amazonaws.com<br><br>sns.eu-west-3.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | sns.eu-south-2.amazonaws.com<br><br>sns.eu-south-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | sns.eu-north-1.amazonaws.com<br><br>sns.eu-north-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | sns.eu-central-2.amazonaws.com<br><br>sns.eu-central-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Israel (Tel Aviv) | il-central-1 | sns.il-central-1.amazonaws.com<br><br>sns.il-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Mexico (Central) | mx-central-1 | sns.mx-central-1.amazonaws.com<br><br>sns.mx-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | sns.me-south-1.amazonaws.com<br><br>sns.me-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | sns.me-central-1.amazonaws.com<br><br>sns.me-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | sns.sa-east-1.amazonaws.com<br><br>sns.sa-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | sns.us-gov-east-1.amazonaws.com<br><br>sns.us-gov-east-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | sns.us-gov-west-1.amazonaws.com<br><br>sns.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

## Service quotas

The following quotas determine how many Amazon SNS resources you can create in your AWS account, and they determine the rate at which you can issue Amazon SNS API requests.

### Amazon SNS resource

To request an increase, submit an SNS quota increase request.

| Resource | Default |
| --- | --- |
| Topics | • Standard: 100,000 per account<br>• FIFO: 1,000 per account |
| Subscriptions | • Standard: 12,500,000 per topic<br>For Firehose delivery streams, 5 per topic, per subscription owner<br>• FIFO: 100 per topic |
| Pending subscriptions | 5,000 per account |
| Account spend threshold for SMS | 1.00 USD per account |
| Delivery rate for promotional SMS messages | 20 messages per second |
| Delivery rate for transactional SMS messages | 20 messages per second |
| Delivery rate for email messages | 10 messages per second per email or email-json subscript ion. This is a hard limit and can't be increased. |
| Maximum number of messages in PublishBatchRequest | 10 PublishBatchReques tEntries |
| Subscription filter policies | • 200 filter policies per topic<br>• 10,000 filter policies per AWS account |
| Message header | The maximum header size is 16,384 bytes (16 KiB). |
| Message size | The maximum message size is 262,144 bytes (256 KiB). To |

| Resource | Default |
| --- | --- |
| | publish messages larger than 256 KiB, you can check the Amazon SNS Extended Client Libraries. The maximum payload size is 2 GB. |

## Amazon SNS API throttling

The following quotas throttle the rate at which you can issue Amazon SNS API requests.

### Hard

The following quotas cannot be increased.

| API | Transactions per second |
| --- | --- |
| AddPermission | 10 |
| CheckIfPhoneNumberIsOptedOut | 50 |
| CreateSMSSandboxPhoneNumber | 1 |
| DeleteSMSSandboxPhoneNumber | 1 |
| GetSMSAttributes | 20 |
| GetSMSSandboxAccountStatus | 10 |
| ListEndpointsByPlatformApplication | 30 |
| ListOriginationNumbers | 1 |
| ListPhoneNumbersOptedOut | 10 |
| ListPlatformApplications | 15 |
| ListSMSSandboxPhoneNumbers | 1 |

| API | Transactions per second |
| --- | --- |
| ListSubscriptions | 30 |
| ListSubscriptionsByTopic | 30 |
| ListTagsForResource | 10 |
| ListTopics | 30 |
| OptInPhoneNumber | 20 |
| RemovePermission | 10 |
| SetSMSAttributes | 1 |
| Subscribe | 100 |
| TagResource | 10 |
| Unsubscribe | 100 |
| UntagResource | 10 |
| VerifySMSSandboxPhoneNumber | 1 |

**Soft**

The following quotas vary by AWS Region. The messages per second quota is based on the number of messages published to an Amazon SNS region, per account, combining `Publish` and `PublishBatch` API requests including both standard and FIFO topic types. To request a soft limit increase, submit an SNS quota increase request.

For example, if your regional quota is 30,000 messages per second, per account, there are a few ways this quota can be reached:

- Using the `Publish` action at a rate of 30,000 API requests per second to publish 30,000 messages (one message per API request).
- Using the `PublishBatch` action at a rate of 3,000 API requests per second to publish 30,000 messages (10 messages per batch API request).

- Using the `Publish` action at a rate of 10,000 API requests per second to publish 10,000 messages (one message per API request) and the `PublishBatch` action at a rate of 2,000 API requests per second to publish 20,000 messages (10 messages per batch API request) for a total of 30,000 messages published per second.

**Publish API throttling per account**

| AWS Regions | Standard topics | FIFO topics* |
|---|---|---|
| US East (N. Virginia) Region | 30,000 messages per second | 30,000 messages per second |
| US West (Oregon) Region<br>Europe (Ireland) Region | 9,000 messages per second | 9,000 messages per second |
| US East (Ohio) Region<br>US West (N. California) Region<br>Asia Pacific (Mumbai) Region<br>Asia Pacific (Seoul) Region<br>Asia Pacific (Singapore) Region<br>Asia Pacific (Sydney) Region<br>Asia Pacific (Tokyo) Region<br>Europe (Frankfurt) Region | 1,500 messages per second | 3,000 messages per second |
| All other supported Regions | 300 messages per second | 3,000 messages per second |

*Amazon SNS FIFO per topic limits exceed the default per account limit when `FifoThroughputScope` is set to `MessageGroup`. Amazon SNS FIFO topics have a maximum per message group limit of 300 messages per second, and Amazon SNS FIFO topics per topic throughput defaults to a limit of 3,000 messages per second or 20MB per second, which ever comes first, when `FifoThroughputScope` is set to `Topic`. Amazon SNS FIFO topics can

experience reduced throughput within a message group for cross Regional deliveries due to the added latency between Regions, and the need to maintain the strict order of messages.

**Other API throttling**

| APIs | AWS Regions | Transactions per second |
|---|---|---|
| ConfirmSubscription | US East (N. Virginia) Region | 3,000 |
| CreatePlatformApplication<br>CreatePlatformEndpoint | US West (Oregon) Region<br>Europe (Ireland) Region | 900 |
| CreateTopic<br>DeleteEndpoint<br>DeletePlatformApplication<br>DeleteTopic<br>GetEndpointAttributes<br>GetDataProtectionPolicy<br>GetPlatformApplicationAttributes<br>GetSubscriptionAttributes | US East (Ohio) Region<br>US West (N. California) Region<br>Asia Pacific (Mumbai) Region<br>Asia Pacific (Seoul) Region<br>Asia Pacific (Singapore) Region<br>Asia Pacific (Sydney) Region<br>Asia Pacific (Tokyo) Region<br>Europe (Frankfurt) Region | 150 |
| GetTopicAttributes<br>SetEndpointAttributes<br>SetPlatformApplicationAttributes<br>SetSubscriptionAttributes<br>SetTopicAttributes | Africa (Cape Town) Region<br>Asia Pacific (Hong Kong) Region<br>Asia Pacific (Hyderabad)<br>Asia Pacific (Osaka) Region<br>Canada (Central) Region<br>China (Beijing) Region | 30 |

| APIs | AWS Regions | Transactions per second |
| --- | --- | --- |
| | China (Ningxia) Region<br>Europe (London) Region<br>Europe (Milan) Region<br>Europe (Paris) Region<br>Europe (Spain)<br>Europe (Stockholm) Region<br>Israel (Tel Aviv) Region<br>Middle East (Bahrain) Region<br>South America (São Paulo) Region | |
| PutDataProtectionPolicy | All Commercial Regions | 1 |

### Message Archiving and Replay

| Policy | AWS Regions | Standard topics | FIFO topics |
| --- | --- | --- | --- |
| ArchivePolicy | All Commercial Regions | N/A | Yes |
| | AWS GovCloud (US) Regions | N/A | Yes |
| ReplayPolicy | All Commercial Regions | N/A | Yes |
| | AWS GovCloud (US) Regions | N/A | Yes |

# Amazon Simple Queue Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, IPv6 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Amazon SQS

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | sqs.us-east-2.amazonaws.com | HTTP and HTTPS |
| | | sqs-fips.us-east-2.api.aws | HTTPS |
| | | sqs-fips.us-east-2.amazonaws.com | HTTPS |
| | | sqs.us-east-2.api.aws | HTTP and HTTPS |
| US East (N. Virginia) | us-east-1 | sqs.us-east-1.amazonaws.com | HTTP and HTTPS |
| | | sqs-fips.us-east-1.api.aws | HTTPS |
| | | sqs-fips.us-east-1.amazonaws.com | HTTPS |
| | | sqs.us-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | | HTTP and HTTPS | |
| US West (N. California) | us-west-1 | sqs.us-west-1.amazonaws.com | HTTP and HTTPS | |
| | | sqs-fips.us-west-1.api.aws | HTTPS | |
| | | sqs-fips.us-west-1.amazonaws.com | HTTPS | |
| | | sqs.us-west-1.api.aws | HTTP and HTTPS | |
| US West (Oregon) | us-west-2 | sqs.us-west-2.amazonaws.com | HTTP and HTTPS | |
| | | sqs-fips.us-west-2.api.aws | HTTPS | |
| | | sqs-fips.us-west-2.amazonaws.com | HTTPS | |
| | | sqs.us-west-2.api.aws | HTTP and HTTPS | |
| Africa (Cape Town) | af-south-1 | sqs.af-south-1.amazonaws.com | HTTP and HTTPS | |
| | | sqs.af-south-1.api.aws | HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | sqs.ap-east-1.amazonaws.com<br><br>sqs.ap-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | sqs.ap-south-2.amazonaws.com<br><br>sqs.ap-south-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | sqs.ap-southeast-3.amazonaws.com<br><br>sqs.ap-southeast-3.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | sqs.ap-southeast-5.amazonaws.com<br><br>sqs.ap-southeast-5.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | sqs.ap-southeast-4.amazonaws.com<br>sqs.ap-southeast-4.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | sqs.ap-south-1.amazonaws.com<br>sqs.ap-south-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | sqs.ap-southeast-6.amazonaws.com<br>sqs.ap-southeast-6.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | sqs.ap-northeast-3.amazonaws.com<br>sqs.ap-northeast-3.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | sqs.ap-northeast-2.amazonaws.com<br>sqs.ap-northeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | sqs.ap-southeast-1.amazonaws.com<br>sqs.ap-southeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | sqs.ap-southeast-2.amazonaws.com<br>sqs.ap-southeast-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | sqs.ap-east-2.amazonaws.com<br>sqs.ap-east-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | sqs.ap-southeast-7.amazonaws.com<br>sqs.ap-southeast-7.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | sqs.ap-northeast-1.amazonaws.com<br>sqs.ap-northeast-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Canada (Central) | ca-central-1 | sqs.ca-central-1.amazonaws.com<br>sqs-fips.ca-central-1.api.aws<br>sqs-fips.ca-central-1.amazonaws.com<br>sqs.ca-central-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | sqs.ca-west-1.amazonaws.com<br>sqs-fips.ca-west-1.api.aws<br>sqs-fips.ca-west-1.amazonaws.com<br>sqs.ca-west-1.api.aws | HTTP and HTTPS<br>HTTPS<br>HTTPS<br>HTTP and HTTPS | |
| Europe (Frankfurt) | eu-central-1 | sqs.eu-central-1.amazonaws.com<br>sqs.eu-central-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (Ireland) | eu-west-1 | sqs.eu-west-1.amazonaws.com<br>sqs.eu-west-1.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |
| Europe (London) | eu-west-2 | sqs.eu-west-2.amazonaws.com<br>sqs.eu-west-2.api.aws | HTTP and HTTPS<br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | sqs.eu-south-1.amazonaws.com<br><br>sqs.eu-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Paris) | eu-west-3 | sqs.eu-west-3.amazonaws.com<br><br>sqs.eu-west-3.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Spain) | eu-south-2 | sqs.eu-south-2.amazonaws.com<br><br>sqs.eu-south-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Europe (Stockholm) | eu-north-1 | sqs.eu-north-1.amazonaws.com<br><br>sqs.eu-north-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Zurich) | eu-central-2 | sqs.eu-central-2.amazonaws.com<br><br>sqs.eu-central-2.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Israel (Tel Aviv) | il-central-1 | sqs.il-central-1.amazonaws.com<br><br>sqs.il-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Mexico (Central) | mx-central-1 | sqs.mx-central-1.amazonaws.com<br><br>sqs.mx-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| Middle East (Bahrain) | me-south-1 | sqs.me-south-1.amazonaws.com<br><br>sqs.me-south-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Middle East (UAE) | me-central-1 | sqs.me-central-1.amazonaws.com<br><br>sqs.me-central-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| South America (São Paulo) | sa-east-1 | sqs.sa-east-1.amazonaws.com<br><br>sqs.sa-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | sqs.us-gov-east-1.amazonaws.com<br><br>sqs.us-gov-east-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | sqs.us-gov-west-1.amazonaws.com<br><br>sqs.us-gov-west-1.api.aws | HTTP and HTTPS<br><br>HTTP and HTTPS | |

## Legacy endpoints

If you use the AWS CLI or SDK for Python, you can use the following legacy endpoints.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | us-east-2.queue.amazonaws.com | HTTP and HTTPS |
| US East (N. Virginia) | us-east-1 | queue.amazonaws.com | HTTP and HTTPS |
| US West (N. California) | us-west-1 | us-west-1.queue.amazonaws.com | HTTP and HTTPS |
| US West (Oregon) | us-west-2 | us-west-2.queue.amazonaws.com | HTTP and HTTPS |
| Africa (Cape Town) | af-south-1 | af-south-1.queue.amazonaws.com | HTTP |
| Asia Pacific (Mumbai) | ap-south-1 | ap-south-1.queue.amazonaws.com | HTTP and HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | ap-northeast-3.queue.amazonaws.com | HTTP and HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ap-northeast-2.queue.amazonaws.com | HTTP and HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ap-southeast-1.queue.amazonaws.com | HTTP and HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ap-southeast-2.queue.amazonaws.com | HTTP and HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ap-northeast-1.queue.amazonaws.com | HTTP and HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | ca-central-1.queue.amazonaws.com | HTTP and HTTPS |
| China (Beijing) | cn-north-1 | cn-north-1.queue.amazonaws.com | HTTP and HTTPS |
| China (Ningxia) | cn-northwest-1 | cn-northwest-1.queue.amazonaws.com | HTTP and HTTPS |
| Europe (Frankfurt) | eu-central-1 | eu-central-1.queue.amazonaws.com | HTTP and HTTPS |
| Europe (Ireland) | eu-west-1 | eu-west-1.queue.amazonaws.com | HTTP and HTTPS |
| Europe (London) | eu-west-2 | eu-west-2.queue.amazonaws.com | HTTP and HTTPS |
| Europe (Paris) | eu-west-3 | eu-west-3.queue.amazonaws.com | HTTP and HTTPS |
| Europe (Stockholm) | eu-north-1 | eu-north-1.queue.amazonaws.com | HTTP and HTTPS |
| South America (São Paulo) | sa-east-1 | sa-east-1.queue.amazonaws.com | HTTP and HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Actions per Queue Policy | Each supported Region: 7 | No | The number of actions in a queue policy. |
| Attributes per Message | Each supported Region: 10 | No | The number of attributes added to a message. |
| Batched Message ID Length | Each supported Region: 80 | No | The length of a batched message ID. |
| Batched Message Throughput for FIFO Queues | Each supported Region: 3,000 | No | The number of batched transactions per second (TPS) for FIFO queues. |
| Batched Message High Throughput for FIFO Queues | Supported Regions (US East (N. Virginia) , US West (Oregon), and Europe (Ireland)): 700,000<br><br>Europe (Frankfurt) and US East (Ohio): 180,000<br><br>Asia Pacific (Mumbai), Asia Pacific (Singapor e), Asia Pacific (Sydney) and Asia Pacific (Tokyo): 90,000 | No | The number of batched transactions per second (TPS) for FIFO queues. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | Europe (London) and South America (São Paulo):45,000<br><br>All other supported Regions: 24,000 | | |
| Conditions per Queue Policy | Each supported Region: 10 | No | The number of conditions in a queue policy. |
| In-Flight Messages per FIFO Queue | Each supported Region: 120,000 | Contact AWS Support | The number of in-flight messages in a FIFO queue. |
| In-Flight Messages per Standard Queue | Each supported Region: 120,000 | Yes | The number of in-flight messages in a standard queue. |
| Message Invisibility Period | Each supported Region: 0 Seconds | No | The length of time, in seconds, for which Amazon SQS retains a message if it isn't deleted. The maximum is 14 days (1,209,600 seconds). |
| Message Size | Each supported Region: 256 Kilobytes | No | The size of a message, in kilobytes. |
| Message Size in Amazon S3 Bucket | Each supported Region: 2 Gigabytes | No | The size of a message, in gigabytes , in an Amazon S3 bucket. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Messages per Batch | Each supported Region: 10 | No | The number of messages in a message batch. |
| Principals per Queue Policy | Each supported Region: 50 | No | The length of time, in minutes, by which to delay the initial delivery of messages to a queue. |
| Queue Name Length | Each supported Region: 80 | No | The queue name length. |
| Queue Policy Size | Each supported Region: 8,192 Bytes | No | The size, in bytes, of a queue policy. |
| Statements per Queue Policy | Each supported Region: 20 | No | The number of tags added to a queue. |
| UTF-8 Queue Tag Key Length | Each supported Region: 128 | No | The length of a UTF-8 queue tag key. |
| UTF-8 Queue Tag Value Length | Each supported Region: 256 | No | The length of a UTF-8 queue tag value. |
| Unbatched Message Throughput for FIFO Queues | Each supported Region: 300 | No | The number of unbatched transacti ons per second (TPS) for FIFO queues. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Unbatched Message High Throughput for FIFO Queues | Supported Regions (US East (N. Virginia), US West (Oregon), and Europe (Ireland)): 70,000<br><br>US East (Ohio) and Europe (Frankfurt): 18,000<br><br>Supported Regions (Asia Pacific (Mumbai), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo)): 9,000<br><br>Europe (London) and South America (São Paulo): 4,500<br><br>All other supported Regions: 2,400 | No | The number of unbatched transactions per second (TPS) for FIFO queues. |

For more information, see Amazon SQS quotas in the *Amazon Simple Queue Service Developer Guide* and the Limits and restrictions section of the *Amazon SQS FAQs*.

## AWS Security Token Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

> **Note**
>
> AWS recommends using Regional STS endpoints within your applications and avoid using the global (legacy) STS endpoint. Regional STS endpoints reduce latency, build in redundancy, and increase session token validity. For more information about configuring your applications to use the regional STS endpoint, see AWS STS Regionalized endpoints in the *AWS SDKs and Tools Reference Guide*. For more information about the global (legacy) AWS STS endpoint, including how to monitor for use of this endpoint, see How to use Regional AWS STS endpoints in the *AWS Security blog*.

## Service endpoints

Although the global (legacy) AWS STS endpoint `https://sts.amazonaws.com` is highly available, it's hosted in a single AWS Region, US East (N. Virginia), and like other endpoints, it doesn't provide automatic failover to endpoints in other Regions. Most Regional endpoints are active by default, but you must manually enable endpoints for some Regions, such as Asia Pacific (Hong Kong). You can deactivate STS endpoints for any Regions that are enabled by default if you do not intend to use those Regions.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | sts.us-east-2.amazonaws.com | HTTPS | |
| | | sts-fips.us-east-2.api.aws | HTTPS | |
| | | sts.us-east-2.api.aws | HTTPS | |
| | | sts-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | sts.us-east-1.amazonaws.com | HTTPS | |
| | | sts-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | sts-fips.us-east-1.api.aws | HTTPS |
| | | sts.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | sts.us-west-1.amazonaws.com | HTTPS |
| | | sts-fips.us-west-1.amazonaws.com | HTTPS |
| | | sts-fips.us-west-1.api.aws | HTTPS |
| | | sts.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | sts.us-west-2.amazonaws.com | HTTPS |
| | | sts-fips.us-west-2.api.aws | HTTPS |
| | | sts-fips.us-west-2.amazonaws.com | HTTPS |
| | | sts.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | sts.af-south-1.amazonaws.com | HTTPS |
| | | sts.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | sts.ap-east-1.amazonaws.com | HTTPS |
| | | sts.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | sts.ap-south-2.amazonaws.com | HTTPS |
| | | sts.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | sts.ap-southeast-3.amazonaws.com | HTTPS |
| | | sts.ap-southeast-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Malaysia) | ap-southeast-5 | sts.ap-southeast-5.amazonaws.com<br>sts.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | sts.ap-southeast-4.amazonaws.com<br>sts.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | sts.ap-south-1.amazonaws.com<br>sts.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | sts.ap-southeast-6.amazonaws.com<br>sts.ap-southeast-6.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | sts.ap-northeast-3.amazonaws.com<br>sts.ap-northeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | sts.ap-northeast-2.amazonaws.com<br>sts.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | sts.ap-southeast-1.amazonaws.com<br>sts.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | sts.ap-southeast-2.amazonaws.com<br>sts.ap-southeast-2.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Taipei) | ap-east-2 | sts.ap-east-2.amazonaws.com | HTTPS |
| | | sts.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | sts.ap-southeast-7.amazonaws.com | HTTPS |
| | | sts.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | sts.ap-northeast-1.amazonaws.com | HTTPS |
| | | sts.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | sts.ca-central-1.amazonaws.com | HTTPS |
| | | sts-fips.ca-central-1.api.aws | HTTPS |
| | | sts.ca-central-1.api.aws | HTTPS |
| | | sts-fips.ca-central-1.amazonaws.com | HTTPS |
| Canada West (Calgary) | ca-west-1 | sts.ca-west-1.amazonaws.com | HTTPS |
| | | sts-fips.ca-west-1.amazonaws.com | HTTPS |
| | | sts-fips.ca-west-1.api.aws | HTTPS |
| | | sts.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | sts.eu-central-1.amazonaws.com | HTTPS |
| | | sts.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | sts.eu-west-1.amazonaws.com | HTTPS |
| | | sts.eu-west-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (London) | eu-west-2 | sts.eu-west-2.amazonaws.com | HTTPS |
| | | sts.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | sts.eu-south-1.amazonaws.com | HTTPS |
| | | sts.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | sts.eu-west-3.amazonaws.com | HTTPS |
| | | sts.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | sts.eu-south-2.amazonaws.com | HTTPS |
| | | sts.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | sts.eu-north-1.amazonaws.com | HTTPS |
| | | sts.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | sts.eu-central-2.amazonaws.com | HTTPS |
| | | sts.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | sts.il-central-1.amazonaws.com | HTTPS |
| | | sts.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | sts.mx-central-1.amazonaws.com | HTTPS |
| | | sts.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | sts.me-south-1.amazonaws.com | HTTPS |
| | | sts.me-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | sts.me-central-1.amazonaws.com<br>sts.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | sts.sa-east-1.amazonaws.com<br>sts.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | sts.us-gov-east-1.amazonaws.com<br>sts.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | sts.us-gov-west-1.amazonaws.com<br>sts.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

### Service quotas

For information about AWS STS quotas, see IAM and AWS STS quotas in the *IAM User Guide*.

## Amazon Simple Workflow Service endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | swf.us-east-2.amazonaws.com | HTTPS | |
| | | swf-fips.us-east-2.api.aws | HTTPS | |
| | | swf-fips.us-east-2.amazonaws.com | HTTPS | |
| | | swf.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | swf.us-east-1.amazonaws.com | HTTPS | |
| | | swf-fips.us-east-1.api.aws | HTTPS | |
| | | swf-fips.us-east-1.amazonaws.com | HTTPS | |
| | | swf.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | swf.us-west-1.amazonaws.com | HTTPS | |
| | | swf-fips.us-west-1.api.aws | HTTPS | |
| | | swf-fips.us-west-1.amazonaws.com | HTTPS | |
| | | swf.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | swf.us-west-2.amazonaws.com | HTTPS | |
| | | swf-fips.us-west-2.api.aws | HTTPS | |
| | | swf-fips.us-west-2.amazonaws.com | HTTPS | |
| | | swf.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | swf.af-south-1.amazonaws.com | HTTPS | |
| | | swf.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | swf.ap-east-1.amazonaws.com<br>swf.ap-east-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | swf.ap-south-2.amazonaws.com<br>swf.ap-south-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | swf.ap-southeast-3.amazonaws.com<br>swf.ap-southeast-3.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | swf.ap-southeast-5.amazonaws.com<br>swf.ap-southeast-5.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | swf.ap-southeast-4.amazonaws.com<br>swf.ap-southeast-4.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | swf.ap-south-1.amazonaws.com<br>swf.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | swf.ap-southeast-6.amazonaws.com<br>swf.ap-southeast-6.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | swf.ap-northeast-3.amazonaws.com<br>swf.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | swf.ap-northeast-2.amazonaws.com<br>swf.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | swf.ap-southeast-1.amazonaws.com<br>swf.ap-southeast-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | swf.ap-southeast-2.amazonaws.com<br>swf.ap-southeast-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | swf.ap-east-2.amazonaws.com<br>swf.ap-east-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | swf.ap-southeast-7.amazonaws.com<br>swf.ap-southeast-7.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | swf.ap-northeast-1.amazonaws.com<br>swf.ap-northeast-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada (Central) | ca-central-1 | swf.ca-central-1.amazonaws.com | HTTPS |
| | | swf-fips.ca-central-1.api.aws | HTTPS |
| | | swf-fips.ca-central-1.amazonaws.com | HTTPS |
| | | swf.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | swf.ca-west-1.amazonaws.com | HTTPS |
| | | swf-fips.ca-west-1.api.aws | HTTPS |
| | | swf-fips.ca-west-1.amazonaws.com | HTTPS |
| | | swf.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | swf.eu-central-1.amazonaws.com | HTTPS |
| | | swf.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | swf.eu-west-1.amazonaws.com | HTTPS |
| | | swf.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | swf.eu-west-2.amazonaws.com | HTTPS |
| | | swf.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | swf.eu-south-1.amazonaws.com | HTTPS |
| | | swf.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | swf.eu-west-3.amazonaws.com | HTTPS |
| | | swf.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | swf.eu-south-2.amazonaws.com | HTTPS |
| | | swf.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | swf.eu-north-1.amazonaws.com | HTTPS |
| | | swf.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | swf.eu-central-2.amazonaws.com | HTTPS |
| | | swf.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | swf.il-central-1.amazonaws.com | HTTPS |
| | | swf.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | swf.mx-central-1.amazonaws.com | HTTPS |
| | | swf.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | swf.me-south-1.amazonaws.com | HTTPS |
| | | swf.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | swf.me-central-1.amazonaws.com | HTTPS |
| | | swf.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | swf.sa-east-1.amazonaws.com | HTTPS |
| | | swf.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | swf.us-gov-east-1.amazonaws.com | HTTPS |
| | | swf-fips.us-gov-east-1.api.aws | HTTPS |
| | | swf-fips.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | swf.us-gov-west-1.amazonaws.com<br>swf-fips.us-gov-west-1.api.aws<br>swf-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| CountClosedWorkflowExecutions throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of CountClosedWorkflowExecutions calls you can burst without being throttled. |
| CountClosedWorkflowExecutions throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of CountClosedWorkflowExecutions calls you can make per second without being throttled. |
| CountOpenWorkflowExecutions throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of CountOpenWorkflowExecutions calls you can burst without being throttled. |
| CountOpenWorkflowExecutions throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of CountOpenWorkflowExecutions calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| CountPendingActivityTasks throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of CountPendingActivityTasks calls you can burst without being throttled. |
| CountPendingActivityTasks throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of CountPendingActivityTasks calls you can make per second without being throttled. |
| CountPendingDecisionTasks throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of CountPendingDecisionTasks calls you can burst without being throttled. |
| CountPendingDecisionTasks throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of CountPendingDecisionTasks calls you can make per second without being throttled. |
| DeleteActivityType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of DeleteActivityType calls you can burst without being throttled. |
| DeleteActivityType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DeleteActivityType calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DeleteWorkflowType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of DeleteWorkflowType calls you can burst without being throttled. |
| DeleteWorkflowType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DeleteWorkflowType calls you can make per second without being throttled. |
| DeprecateActivityType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of DeprecateActivityType calls you can burst without being throttled. |
| DeprecateActivityType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DeprecateActivityType calls you can make per second without being throttled. |
| DeprecateDomain throttle burst limit in transactions per second | Each supported Region: 100 | Yes | The maximum number of DeprecateDomain calls you can burst without being throttled. |
| DeprecateDomain throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DeprecateDomain calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DeprecateWorkflowType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of DeprecateWorkflowType calls you can burst without being throttled. |
| DeprecateWorkflowType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DeprecateWorkflowType calls you can make per second without being throttled. |
| DescribeActivityType throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of DescribeActivityType calls you can burst without being throttled. |
| DescribeActivityType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DescribeActivityType calls you can make per second without being throttled. |
| DescribeDomain throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of DescribeDomain calls you can burst without being throttled. |
| DescribeDomain throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DescribeDomain calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DescribeWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of DescribeWorkflowExecution calls you can burst without being throttled. |
| DescribeWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DescribeWorkflowExecution calls you can make per second without being throttled. |
| DescribeWorkflowType throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of DescribeWorkflowType calls you can burst without being throttled. |
| DescribeWorkflowType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of DescribeWorkflowType calls you can make per second without being throttled. |
| Events in Workflow execution history | Each supported Region: 25,000 | No | The maximum number of events for a given workflow execution. |
| GetWorkflowExecutionHistory throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of GetWorkflowExecutionHistory calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetWorkflowExecutionHistory throttle refill limit in transactions per second | Each supported Region: 60 | Yes | The maximum number of GetWorkflowExecutionHistory calls you can make per second without being throttled. |
| Input or output result data size | Each supported Region: 32,768 | No | This limit affects activity or workflow execution result data, input data when scheduling activity tasks or workflow executions, and input sent with a workflow execution signal. |
| ListActivityTypes throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of ListActivityTypes calls you can burst without being throttled. |
| ListActivityTypes throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of ListActivityTypes calls you can make per second without being throttled. |
| ListClosedWorkflowExecutions throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of ListClosedWorkflowExecutions calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListClosedWorkflowExecutions throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of ListClosedWorkflowExecutions calls you can make per second without being throttled. |
| ListDomains throttle burst limit in transactions per second | Each supported Region: 100 | Yes | The maximum number of ListDomains calls you can burst without being throttled. |
| ListDomains throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of ListDomains calls you can make per second without being throttled. |
| ListOpenWorkflowExecutions throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of ListOpenWorkflowExecutions calls you can burst without being throttled. |
| ListOpenWorkflowExecutions throttle refill limit in transactions per second | Each supported Region: 48 | Yes | The maximum number of ListOpenWorkflowExecutions calls you can make per second without being throttled. |
| ListTagsForResource throttle burst limit in transactions per second | Each supported Region: 50 | Yes | The maximum number of ListTagsForResource calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListTagsForResource throttle refill limit in transactions per second | Each supported Region: 30 | Yes | The maximum number of ListTagsForResource calls you can make per second without being throttled. |
| ListWorkflowTypes throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of ListWorkflowTypes calls you can burst without being throttled. |
| ListWorkflowTypes throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of ListWorkflowTypes calls you can make per second without being throttled. |
| Open activity tasks per workflow execution | Each supported Region: 1,000 | No | This limit includes both activity tasks that have been scheduled and those being processed by workers. |
| Open child workflow executions | Each supported Region: 1,000 | No | The maximum number of open child workflow executions per workflow execution. |
| Open timers per workflow execution | Each supported Region: 1,000 | No | The maximum number of concurrently open timers per workflow execution. |
| Open workflow executions per domain | Each supported Region: 100,000 | Yes | The maximum number of open workflow executions per domain for this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| PollForActivityTask throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of PollForActivityTask calls you can burst without being throttled. |
| PollForActivityTask throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of PollForActivityTask calls you can make per second without being throttled. |
| PollForDecisionTask throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of PollForDecisionTask calls you can burst without being throttled. |
| PollForDecisionTask throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of PollForDecisionTask calls you can make per second without being throttled. |
| Pollers per task list | Each supported Region: 1,000 | No | You can have a maximum of 1,000 pollers which simultaneously poll a particular task list. |
| RecordActivityTaskHeartbeat throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of RecordActivityTaskHeartbeat calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| RecordActivityTaskHeartbeat throttle refill limit in transactions per second | Each supported Region: 160 | Yes | The maximum number of RecordActivityTaskHeartbeat calls you can make per second without being throttled. |
| RegisterActivityType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of RegisterActivityType calls you can burst without being throttled. |
| RegisterActivityType throttle refill limit in transactions per second | Each supported Region: 60 | Yes | The maximum number of RegisterActivityType calls you can make per second without being throttled. |
| RegisterDomain throttle burst limit in transactions per second | Each supported Region: 100 | Yes | The maximum number of RegisterDomain calls you can burst without being throttled. |
| RegisterDomain throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of RegisterDomain calls you can make per second without being throttled. |
| RegisterWorkflowType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of RegisterWorkflowType calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| RegisterWorkflowType throttle refill limit in transactions per second | Each supported Region: 60 | Yes | The maximum number of RegisterWorkflowType calls you can make per second without being throttled. |
| Registered domains | Each supported Region: 100 | Yes | The maximum number of registered domains for this account in the current Region. |
| Request size | Each supported Region: 1 Megabytes | No | The total data size in megabytes per Simple Workflow API request, including the request header and all other associated request data. |
| RequestCancelExternalWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 1,200 | Yes | The maximum number of RequestCancelExternalWorkflowExecution calls you can burst without being throttled. |
| RequestCancelExternalWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 120 | Yes | The maximum number of RequestCancelExternalWorkflowExecution calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| RequestCancelWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of RequestCancelWorkflowExecution calls you can burst without being throttled. |
| RequestCancelWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 30 | Yes | The maximum number of RequestCancelWorkflowExecution calls you can make per second without being throttled. |
| RespondActivityTaskCanceled throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of RespondActivityTaskCanceled calls you can burst without being throttled. |
| RespondActivityTaskCanceled throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of RespondActivityTaskCanceled calls you can make per second without being throttled. |
| RespondActivityTaskCompleted throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of RespondActivityTaskCompleted calls you can burst without being throttled. |
| RespondActivityTaskCompleted throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of RespondActivityTaskCompleted calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| RespondActivityTaskFailed throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of RespondActivityTaskFailed calls you can burst without being throttled. |
| RespondActivityTaskFailed throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of RespondActivityTaskFailed calls you can make per second without being throttled. |
| RespondDecisionTaskCompleted throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of RespondDecisionTaskCompleted calls you can burst without being throttled. |
| RespondDecisionTaskCompleted throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of RespondDecisionTaskCompleted calls you can make per second without being throttled. |
| SWF task retention time in queue | Each supported Region: 1 | No | The maximum time, in years, for a task to stay in queued state (constrained by workflow execution time limit). |
| ScheduleActivityTask throttle burst limit in transactions per second | Each supported Region: 1,000 | Yes | The maximum number of ScheduleActivityTask calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ScheduleActivityTask throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of ScheduleActivityTask calls you can make per second without being throttled. |
| SignalExternalWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 1,200 | Yes | The maximum number of SignalExternalWorkflowExecution calls you can burst without being throttled. |
| SignalExternalWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 120 | Yes | The maximum number of SignalExternalWorkflowExecution calls you can make per second without being throttled. |
| SignalWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of SignalWorkflowExecution calls you can burst without being throttled. |
| SignalWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 30 | Yes | The maximum number of SignalWorkflowExecution calls you can make per second without being throttled. |
| StartChildWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 500 | Yes | The maximum number of StartChildWorkflowExecution calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StartChildWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 12 | Yes | The maximum number of StartChildWorkflowExecution calls you can make per second without being throttled. |
| StartTimer throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of StartTimer calls you can burst without being throttled. |
| StartTimer throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of StartTimer calls you can make per second without being throttled. |
| StartWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of StartWorkflowExecution calls you can burst without being throttled. |
| StartWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of StartWorkflowExecution calls you can make per second without being throttled. |
| TagResource throttle burst limit in transactions per second | Each supported Region: 50 | Yes | The maximum number of TagResource calls you can burst without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| TagResource throttle refill limit in transactions per second | Each supported Region: 30 | Yes | The maximum number of TagResource calls you can make per second without being throttled. |
| Task execution time in year | Each supported Region: 1 | No | The maximum time for a task to stay in execution state (constrained by workflow execution time limit). |
| TerminateWorkflowExecution throttle burst limit in transactions per second | Each supported Region: 2,000 | Yes | The maximum number of TerminateWorkflowExecution calls you can burst without being throttled. |
| TerminateWorkflowExecution throttle refill limit in transactions per second | Each supported Region: 60 | Yes | The maximum number of TerminateWorkflowExecution calls you can make per second without being throttled. |
| UndeprecateActivityType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of UndeprecateActivityType calls you can burst without being throttled. |
| UndeprecateActivityType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of UndeprecateActivityType calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| UndeprecateDomain throttle burst limit in transactions per second | Each supported Region: 100 | Yes | The maximum number of UndeprecateDomain calls you can burst without being throttled. |
| UndeprecateDomain throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of UndeprecateDomain calls you can make per second without being throttled. |
| UndeprecateWorkflowType throttle burst limit in transactions per second | Each supported Region: 200 | Yes | The maximum number of UndeprecateWorkflowType calls you can burst without being throttled. |
| UndeprecateWorkflowType throttle refill limit in transactions per second | Each supported Region: 6 | Yes | The maximum number of UndeprecateWorkflowType calls you can make per second without being throttled. |
| UntagResource throttle burst limit in transactions per second | Each supported Region: 50 | Yes | The maximum number of UntagResource calls you can burst without being throttled. |
| UntagResource throttle refill limit in transactions per second | Each supported Region: 30 | Yes | The maximum number of UntagResource calls you can make per second without being throttled. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Workflow and activity types per domain | Each supported Region: 10,000 | Yes | The maximum number of registered workflow and activity types per domain for this account in the current Region. |
| Workflow execution time | Each supported Region: 1 | No | The maximum amount of time, in years, for which a workflow execution can run. |

For more information, see Amazon SWF Quotas in the *Amazon Simple Workflow Service Developer Guide*.

# AWS Snow Family endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Snow Family devices are available in the following AWS Regions.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | snowball.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | snowball.us-east-2.api.aws | HTTPS | |
| | | snowball-fips.us-east-2.amazonaws.com | HTTPS | |
| | | snowball-fips.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | snowball.us-east-1.amazonaws.com | HTTPS | |
| | | snowball-fips.us-east-1.amazonaws.com | HTTPS | |
| | | snowball.us-east-1.api.aws | HTTPS | |
| | | snowball-fips.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | snowball.us-west-1.amazonaws.com | HTTPS | |
| | | snowball.us-west-1.api.aws | HTTPS | |
| | | snowball-fips.us-west-1.api.aws | HTTPS | |
| | | snowball-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | snowball.us-west-2.amazonaws.com | HTTPS | |
| | | snowball-fips.us-west-2.amazonaws.com | HTTPS | |
| | | snowball.us-west-2.api.aws | HTTPS | |
| | | snowball-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | snowball.af-south-1.amazonaws.com | HTTPS | |
| | | snowball.af-south-1.api.aws | HTTPS | |
| | | snowball-fips.af-south-1.amazonaws.com | HTTPS | |
| | | snowball-fips.af-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Hong Kong) | ap-east-1 | snowball.ap-east-1.amazonaws.com | HTTPS |
| | | snowball.ap-east-1.api.aws | HTTPS |
| | | snowball-fips.ap-east-1.api.aws | HTTPS |
| | | snowball-fips.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | snowball.ap-southeast-3.amazonaws.com | HTTPS |
| | | snowball.ap-southeast-3.api.aws | HTTPS |
| | | snowball-fips.ap-southeast-3.amazonaws.com | HTTPS |
| | | snowball-fips.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | snowball.ap-south-1.amazonaws.com | HTTPS |
| | | snowball.ap-south-1.api.aws | HTTPS |
| | | snowball-fips.ap-south-1.amazonaws.com | HTTPS |
| | | snowball-fips.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | snowball.ap-northeast-3.amazonaws.com | HTTPS |
| | | snowball-fips.ap-northeast-3.api.aws | HTTPS |
| | | snowball-fips.ap-northeast-3.amazonaws.com | HTTPS |
| | | snowball.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | snowball.ap-northeast-2.amazonaws.com | HTTPS |
| | | snowball-fips.ap-northeast-2.amazonaws.com | HTTPS |
| | | snowball.ap-northeast-2.api.aws | HTTPS |
| | | snowball-fips.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | snowball.ap-southeast-1.amazonaws.com | HTTPS | |
| | | snowball.ap-southeast-1.api.aws | HTTPS | |
| | | snowball-fips.ap-southeast-1.api.aws | HTTPS | |
| | | snowball-fips.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | snowball.ap-southeast-2.amazonaws.com | HTTPS | |
| | | snowball-fips.ap-southeast-2.api.aws | HTTPS | |
| | | snowball-fips.ap-southeast-2.amazonaws.com | HTTPS | |
| | | snowball.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | snowball.ap-northeast-1.amazonaws.com | HTTPS | |
| | | snowball-fips.ap-northeast-1.amazonaws.com | HTTPS | |
| | | snowball.ap-northeast-1.api.aws | HTTPS | |
| | | snowball-fips.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | snowball.ca-central-1.amazonaws.com | HTTPS | |
| | | snowball.ca-central-1.api.aws | HTTPS | |
| | | snowball-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | snowball-fips.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | snowball.eu-central-1.amazonaws.com | HTTPS | |
| | | snowball.eu-central-1.api.aws | HTTPS | |
| | | snowball-fips.eu-central-1.api.aws | HTTPS | |
| | | snowball-fips.eu-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | snowball.eu-west-1.amazonaws.com | HTTPS |
| | | snowball-fips.eu-west-1.amazonaws.com | HTTPS |
| | | snowball.eu-west-1.api.aws | HTTPS |
| | | snowball-fips.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | snowball.eu-west-2.amazonaws.com | HTTPS |
| | | snowball.eu-west-2.api.aws | HTTPS |
| | | snowball-fips.eu-west-2.amazonaws.com | HTTPS |
| | | snowball-fips.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | snowball.eu-south-1.amazonaws.com | HTTPS |
| | | snowball-fips.eu-south-1.api.aws | HTTPS |
| | | snowball-fips.eu-south-1.amazonaws.com | HTTPS |
| | | snowball.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | snowball.eu-west-3.amazonaws.com | HTTPS |
| | | snowball-fips.eu-west-3.api.aws | HTTPS |
| | | snowball-fips.eu-west-3.amazonaws.com | HTTPS |
| | | snowball.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | snowball.eu-north-1.amazonaws.com | HTTPS |
| | | snowball.eu-north-1.api.aws | HTTPS |
| | | snowball-fips.eu-north-1.amazonaws.com | HTTPS |
| | | snowball-fips.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | snowball.il-central-1.amazonaws.com | HTTPS | |
| | | snowball-fips.il-central-1.api.aws | HTTPS | |
| | | snowball-fips.il-central-1.amazonaws.com | HTTPS | |
| | | snowball.il-central-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | snowball.me-central-1.amazonaws.com | HTTPS | |
| | | snowball-fips.me-central-1.api.aws | HTTPS | |
| | | snowball-fips.me-central-1.amazonaws.com | HTTPS | |
| | | snowball.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | snowball.sa-east-1.amazonaws.com | HTTPS | |
| | | snowball.sa-east-1.api.aws | HTTPS | |
| | | snowball-fips.sa-east-1.amazonaws.com | HTTPS | |
| | | snowball-fips.sa-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | snowball.us-gov-east-1.amazonaws.com | HTTPS | |
| | | snowball.us-gov-east-1.api.aws | HTTPS | |
| | | snowball-fips.us-gov-east-1.amazonaws.com | HTTPS | |
| | | snowball-fips.us-gov-east-1.api.aws | HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | snowball.us-gov-west-1.amazonaws.com | HTTPS | |
| | | snowball-fips.us-gov-west-1.api.aws | HTTPS | |
| | | snowball-fips.us-gov-west-1.amazonaws.com | HTTPS | |
| | | snowball.us-gov-west-1.api.aws | HTTPS | |

**AWS Snowball Edge is available only in the following AWS Regions:**

- US East (N. Virginia)
- US East (Ohio)
- US West (N. California)
- US West (Oregon)
- Canada (Central)
- South America (São Paulo)
- Europe (Ireland)
- Europe (Frankfurt)
- Europe (London)
- Europe (Paris)
- Asia Pacific (Mumbai)
- Asia Pacific (Sydney)
- Asia Pacific (Singapore)
- Asia Pacific (Tokyo)
- Israel (Tel Aviv)

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Snowball Edge devices | Each supported Region: 1 | Yes | The maximum number of Snowball Edge devices. |
| Snowcone devices | Each supported Region: 1 | Yes | The maximum number of Snowcone devices. |

# AWS Step Functions endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | states.us-east-2.amazonaws.com | HTTPS | |
| | | sync-states.us-east-2.api.aws | HTTPS | |
| | | sync-states-fips.us-east-2.api.aws | HTTPS | |
| | | sync-states-fips.us-east-2.amazonaws.com | HTTPS | |
| | | states-fips.us-east-2.api.aws | HTTPS | |
| | | states-fips.us-east-2.amazonaws.com | HTTPS | |
| | | sync-states.us-east-2.amazonaws.com | HTTPS | |
| | | states.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | states.us-east-1.amazonaws.com | HTTPS | |
| | | sync-states.us-east-1.api.aws | HTTPS | |
| | | sync-states-fips.us-east-1.api.aws | HTTPS | |
| | | sync-states-fips.us-east-1.amazonaws.com | HTTPS | |
| | | states-fips.us-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | states-fips.us-east-1.amazonaws.com | HTTPS | |
| | | sync-states.us-east-1.amazonaws.com | HTTPS | |
| | | states.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | states.us-west-1.amazonaws.com | HTTPS | |
| | | sync-states.us-west-1.api.aws | HTTPS | |
| | | sync-states-fips.us-west-1.api.aws | HTTPS | |
| | | sync-states-fips.us-west-1.amazonaws.com | HTTPS | |
| | | states-fips.us-west-1.api.aws | HTTPS | |
| | | states-fips.us-west-1.amazonaws.com | HTTPS | |
| | | sync-states.us-west-1.amazonaws.com | HTTPS | |
| | | states.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | states.us-west-2.amazonaws.com | HTTPS | |
| | | sync-states.us-west-2.api.aws | HTTPS | |
| | | sync-states-fips.us-west-2.api.aws | HTTPS | |
| | | sync-states-fips.us-west-2.amazonaws.com | HTTPS | |
| | | states-fips.us-west-2.api.aws | HTTPS | |
| | | states-fips.us-west-2.amazonaws.com | HTTPS | |
| | | sync-states.us-west-2.amazonaws.com | HTTPS | |
| | | states.us-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Africa (Cape Town) | af-south-1 | states.af-south-1.amazonaws.com | HTTPS | |
| | | sync-states.af-south-1.api.aws | HTTPS | |
| | | sync-states.af-south-1.amazonaws.com | HTTPS | |
| | | states.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | states.ap-east-1.amazonaws.com | HTTPS | |
| | | sync-states.ap-east-1.api.aws | HTTPS | |
| | | sync-states.ap-east-1.amazonaws.com | HTTPS | |
| | | states.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Hyderabad) | ap-south-2 | states.ap-south-2.amazonaws.com | HTTPS | |
| | | sync-states.ap-south-2.api.aws | HTTPS | |
| | | sync-states.ap-south-2.amazonaws.com | HTTPS | |
| | | states.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | states.ap-southeast-3.amazonaws.com | HTTPS | |
| | | sync-states.ap-southeast-3.api.aws | HTTPS | |
| | | sync-states.ap-southeast-3.amazonaws.com | HTTPS | |
| | | states.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | states.ap-southeast-5.amazonaws.com | HTTPS | |
| | | sync-states.ap-southeast-5.api.aws | HTTPS | |
| | | sync-states.ap-southeast-5.amazonaws.com | HTTPS | |
| | | states.ap-southeast-5.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Melbourne) | ap-southeast-4 | states.ap-southeast-4.amazonaws.com | HTTPS |
| | | sync-states.ap-southeast-4.api.aws | HTTPS |
| | | sync-states.ap-southeast-4.amazonaws.com | HTTPS |
| | | states.ap-southeast-4.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | states.ap-south-1.amazonaws.com | HTTPS |
| | | sync-states.ap-south-1.api.aws | HTTPS |
| | | sync-states.ap-south-1.amazonaws.com | HTTPS |
| | | states.ap-south-1.api.aws | HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | states.ap-southeast-6.amazonaws.com | HTTPS |
| | | sync-states.ap-southeast-6.api.aws | HTTPS |
| | | sync-states.ap-southeast-6.amazonaws.com | HTTPS |
| | | states.ap-southeast-6.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | states.ap-northeast-3.amazonaws.com | HTTPS |
| | | sync-states.ap-northeast-3.api.aws | HTTPS |
| | | sync-states.ap-northeast-3.amazonaws.com | HTTPS |
| | | states.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | states.ap-northeast-2.amazonaws.com | HTTPS |
| | | sync-states.ap-northeast-2.api.aws | HTTPS |
| | | sync-states.ap-northeast-2.amazonaws.com | HTTPS |
| | | states.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | states.ap-southeast-1.amazonaws.com | HTTPS | |
| | | sync-states.ap-southeast-1.api.aws | HTTPS | |
| | | sync-states.ap-southeast-1.amazonaws.com | HTTPS | |
| | | states.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | states.ap-southeast-2.amazonaws.com | HTTPS | |
| | | sync-states.ap-southeast-2.api.aws | HTTPS | |
| | | sync-states.ap-southeast-2.amazonaws.com | HTTPS | |
| | | states.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | states.ap-east-2.amazonaws.com | HTTPS | |
| | | sync-states.ap-east-2.api.aws | HTTPS | |
| | | sync-states.ap-east-2.amazonaws.com | HTTPS | |
| | | states.ap-east-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | states.ap-southeast-7.amazonaws.com | HTTPS | |
| | | sync-states.ap-southeast-7.api.aws | HTTPS | |
| | | sync-states.ap-southeast-7.amazonaws.com | HTTPS | |
| | | states.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | states.ap-northeast-1.amazonaws.com | HTTPS | |
| | | sync-states.ap-northeast-1.api.aws | HTTPS | |
| | | sync-states.ap-northeast-1.amazonaws.com | HTTPS | |
| | | states.ap-northeast-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | states.ca-central-1.amazonaws.com | HTTPS | |
| | | sync-states.ca-central-1.api.aws | HTTPS | |
| | | sync-states-fips.ca-central-1.api.aws | HTTPS | |
| | | sync-states-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | states-fips.ca-central-1.api.aws | HTTPS | |
| | | sync-states.ca-central-1.amazonaws.com | HTTPS | |
| | | states-fips.ca-central-1.amazonaws.com | HTTPS | |
| | | states.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | states.ca-west-1.amazonaws.com | HTTPS | |
| | | sync-states.ca-west-1.api.aws | HTTPS | |
| | | sync-states-fips.ca-west-1.api.aws | HTTPS | |
| | | sync-states-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | states-fips.ca-west-1.api.aws | HTTPS | |
| | | sync-states.ca-west-1.amazonaws.com | HTTPS | |
| | | states-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | states.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | states.eu-central-1.amazonaws.com | HTTPS | |
| | | sync-states.eu-central-1.api.aws | HTTPS | |
| | | sync-states.eu-central-1.amazonaws.com | HTTPS | |
| | | states.eu-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | states.eu-west-1.amazonaws.com | HTTPS | |
| | | sync-states.eu-west-1.api.aws | HTTPS | |
| | | sync-states.eu-west-1.amazonaws.com | HTTPS | |
| | | states.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | states.eu-west-2.amazonaws.com | HTTPS | |
| | | sync-states.eu-west-2.api.aws | HTTPS | |
| | | sync-states.eu-west-2.amazonaws.com | HTTPS | |
| | | states.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | states.eu-south-1.amazonaws.com | HTTPS | |
| | | sync-states.eu-south-1.api.aws | HTTPS | |
| | | sync-states.eu-south-1.amazonaws.com | HTTPS | |
| | | states.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | states.eu-west-3.amazonaws.com | HTTPS | |
| | | sync-states.eu-west-3.api.aws | HTTPS | |
| | | sync-states.eu-west-3.amazonaws.com | HTTPS | |
| | | states.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | states.eu-south-2.amazonaws.com | HTTPS | |
| | | sync-states.eu-south-2.api.aws | HTTPS | |
| | | sync-states.eu-south-2.amazonaws.com | HTTPS | |
| | | states.eu-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | states.eu-north-1.amazonaws.com | HTTPS |
| | | sync-states.eu-north-1.api.aws | HTTPS |
| | | sync-states.eu-north-1.amazonaws.com | HTTPS |
| | | states.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | states.eu-central-2.amazonaws.com | HTTPS |
| | | sync-states.eu-central-2.api.aws | HTTPS |
| | | sync-states.eu-central-2.amazonaws.com | HTTPS |
| | | states.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | states.il-central-1.amazonaws.com | HTTPS |
| | | sync-states.il-central-1.api.aws | HTTPS |
| | | sync-states.il-central-1.amazonaws.com | HTTPS |
| | | states.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | states.mx-central-1.amazonaws.com | HTTPS |
| | | sync-states.mx-central-1.api.aws | HTTPS |
| | | sync-states.mx-central-1.amazonaws.com | HTTPS |
| | | states.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | states.me-south-1.amazonaws.com | HTTPS |
| | | sync-states.me-south-1.api.aws | HTTPS |
| | | sync-states.me-south-1.amazonaws.com | HTTPS |
| | | states.me-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (UAE) | me-central-1 | states.me-central-1.amazonaws.com | HTTPS |
| | | sync-states.me-central-1.api.aws | HTTPS |
| | | sync-states.me-central-1.amazonaws.com | HTTPS |
| | | states.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | states.sa-east-1.amazonaws.com | HTTPS |
| | | sync-states.sa-east-1.api.aws | HTTPS |
| | | sync-states.sa-east-1.amazonaws.com | HTTPS |
| | | states.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | states.us-gov-east-1.amazonaws.com | HTTPS |
| | | sync-states.us-gov-east-1.api.aws | HTTPS |
| | | sync-states-fips.us-gov-east-1.api.aws | HTTPS |
| | | sync-states-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | states-fips.us-gov-east-1.api.aws | HTTPS |
| | | states-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | sync-states.us-gov-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | states.us-gov-west-1.amazonaws.com | HTTPS | |
| | | sync-states.us-gov-west-1.api.aws | HTTPS | |
| | | sync-states-fips.us-gov-west-1.api.aws | HTTPS | |
| | | sync-states.us-gov-west-1.amazonaws.com | HTTPS | |
| | | states-fips.us-gov-west-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Open Map Runs | Each supported Region: 1,000 | No | The maximum number of open Map Runs allowed at a time in this account. |
| Activity pollers per ARN | Each supported Region: 1,000 | No | The maximum number of workers concurrently polling an activity per ARN. |
| CreateActivity throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of CreateActivity calls that you can make per second. |
| CreateActivity throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for CreateActivity calls. |
| CreateStateMachine throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of CreateStateMachine |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | calls that you can make per second. |
| CreateStateMachine throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of CreateStateMachine calls. |
| CreateStateMachineAlias throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of CreateStateMachineAlias calls that you can make per second. |
| CreateStateMachineAlias throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for CreateStateMachineAlias calls. |
| DeleteActivity throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of DeleteActivity calls that you can make per second. |
| DeleteActivity throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of DeleteActivity calls. |
| DeleteStateMachine throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of DeleteStateMachine calls that you can make per second. |
| DeleteStateMachine throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of DeleteStateMachine calls. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DeleteStateMachineAlias throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of DeleteStateMachineAlias calls that you can make per second. |
| DeleteStateMachineAlias throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for DeleteStateMachineAlias calls. |
| DeleteStateMachineVersion throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of DeleteStateMachineVersion calls that you can make per second. |
| DeleteStateMachineVersion throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for DeleteStateMachineVersion calls. |
| DescribeActivity throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of DescribeActivity calls that you can make per second. |
| DescribeActivity throttle token refill rate per second | Each supported Region: 20 | Yes | The token refill rate per second of DescribeActivity calls. |
| DescribeExecution throttle token bucket size | Each supported Region: 300 | Yes | The maximum number of DescribeExecution calls that you can make per second. |
| DescribeExecution throttle token refill rate per second | Each supported Region: 50 | Yes | The token refill rate per second of DescribeExecution calls. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DescribeMapRun throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of DescribeMapRun calls that you can make per second. |
| DescribeMapRun throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for DescribeMapRun calls. |
| DescribeStateMachine throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of DescribeStateMachine calls that you can make per second. |
| DescribeStateMachine throttle token refill rate per second | Each supported Region: 20 | Yes | The token refill rate per second of DescribeStateMachine calls. |
| DescribeStateMachineAlias throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of DescribeStateMachineAlias calls that you can make per second. |
| DescribeStateMachineAlias throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for DescribeStateMachineAlias calls. |
| DescribeStateMachineForExecution throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of DescribeStateMachineForExecution calls that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| DescribeStateMachineForExecution throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of DescribeStateMachineForExecution calls. |
| Execution history retention time | Each supported Region: 90 | Yes | The maximum time (in days) for which the execution history is stored after an execution completes. |
| Execution history size | Each supported Region: 25,000 | No | The maximum number of events in a state machines execution history. |
| Execution idle time | Each supported Region: 1 | No | The maximum amount of time (in years) a state machine execution can remain idle. |
| Execution time | Each supported Region: 1 | No | The maximum amount of time (in years) for which an execution can run. |
| GetActivityTask throttle token bucket size | us-east-1: 3,000<br>us-west-2: 3,000<br>eu-west-1: 3,000<br>Each of the other supported Regions: 1,500 | Yes | The maximum number of GetActivityTask calls that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetActivityTask throttle token refill rate per second | us-east-1: 500<br><br>us-west-2: 500<br><br>eu-west-1: 500<br><br>Each of the other supported Regions: 300 | Yes | The token refill rate per second of GetActivityTask calls. |
| GetExecutionHistory throttle token bucket size | Each supported Region: 400 | Yes | The maximum number of GetExecutionHistory calls that you can make per second. |
| GetExecutionHistory throttle token refill rate per second | Each supported Region: 20 | Yes | The token refill rate per second of GetExecutionHistory calls. |
| Input file data size for a Map Run | Each supported Region: 10 Gigabytes | No | Input file data size for a Map Run |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Input or result data size in task state or execution | Each supported Region: 262,144 Bytes | No | The maximum size of input or output data (in bytes) as a UTF-8 encoded string. This quota affects tasks (activity, Lambda function, or integrated service), state or execution output, and input data when scheduling a task, entering a state, or starting an execution. |
| ListActivities throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ListActivities calls that you can make per second. |
| ListActivities throttle token refill rate per second | us-east-1: 10<br>us-west-2: 10<br>eu-west-1: 10<br>Each of the other supported Regions: 5 | Yes | The token refill rate per second of ListActivities calls. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListExecutions throttle token bucket size | us-east-1: 200<br>us-west-2: 200<br>eu-west-1: 200<br>Each of the other supported Regions: 100 | Yes | The maximum number of ListExecutions calls that you can make per second. |
| ListExecutions throttle token refill rate per second | us-east-1: 5<br>us-west-2: 5<br>eu-west-1: 5<br>Each of the other supported Regions: 2 | Yes | The token refill rate per second of ListExecutions calls. |
| ListMapRuns throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ListMapRuns calls that you can make per second. |
| ListMapRuns throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for ListMapRuns calls. |
| ListStateMachineAliases throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ListStateMachineAliases calls that you can make per second. |
| ListStateMachineAliases throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for ListState MachineAliases calls. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| ListStateMachineVersions throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ListStateMachineVersions calls that you can make per second. |
| ListStateMachineVersions throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for ListStateMachineVersions calls. |
| ListStateMachines throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ListStateMachine calls that you can make per second. |
| ListStateMachines throttle token refill rate per second | Each supported Region: 5 | Yes | The token refill rate per second of ListStateMachines calls. |
| ListTagsForResource throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ListTagsForResource calls that you can make per second. |
| ListTagsForResource throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of ListTagsForResource calls. |
| Maximum Concurrency | Each supported Region: 10,000 | No | The maximum number of parallel Map Run child executions |
| Maximum Number of Items Read per Map Run | Each supported Region: 100,000,000 | No | Maximum Number of Items Read per Map Run |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum redrives of a Map Run | Each supported Region: 1,000 | No | Maximum redrives of a Map Run |
| Maximum size of CSV headers | Each supported Region: 10 Kilobytes | No | Maximum size of CSV headers |
| Open executions | Each supported Region: 1,000,000 | Yes | The maximum number of open executions per account in the current Region. |
| PublishStateMachineVersion throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of PublishStateMachineVersion calls that you can make per second. |
| PublishStateMachineVersion throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for PublishStateMachineVersion calls. |
| Published state machine versions per state machine | Each supported Region: 1,000 | Yes | The maximum number of versions that can be published for a state machine. |
| RedriveExecution throttle token bucket size | Each supported Region: 800 | Yes | The maximum number of RedriveExecution calls that you can make per second. |
| RedriveExecution throttle token refill rate per second | Each supported Region: 150 | Yes | The token refill rate per second for RedriveExecution calls. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Registered activities | Each supported Region: 100,000 | Yes | The maximum number of activities that you can create in this account in the current Region. |
| Registered state machines | Each supported Region: 100,000 | Yes | The maximum number of state machines that you can create in this account in the current Region. |
| Request size | Each supported Region: 1 Megabytes | No | The total data size in Megabytes per Step Functions API request, including the request header and all other associated request data. |
| Resource name length | Each supported Region: 80 | No | The maximum allowed characters for names of state machines, executions, and activity resource types. |
| SendTaskFailure throttle token bucket size | us-east-1: 3,000<br>us-west-2: 3,000<br>eu-west-1: 3,000<br>Each of the other supported Regions: 1,500 | Yes | The maximum number of SendTaskFailure calls that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SendTaskFailure throttle token refill rate per second | us-east-1: 500<br>us-west-2: 500<br>eu-west-1: 500<br>Each of the other supported Regions: 300 | Yes | The token refill rate per second of SendTaskFailure calls. |
| SendTaskHeartbeat throttle token bucket size | us-east-1: 3,000<br>us-west-2: 3,000<br>eu-west-1: 3,000<br>Each of the other supported Regions: 1,500 | Yes | The maximum number of SendTaskHeartbeat calls that you can make per second. |
| SendTaskHeartbeat throttle token refill rate per second | us-east-1: 500<br>us-west-2: 500<br>eu-west-1: 500<br>Each of the other supported Regions: 300 | Yes | The token refill rate per second of SendTaskHeartbeat calls. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| SendTaskSuccess throttle token bucket size | us-east-1: 3,000<br>us-west-2: 3,000<br>eu-west-1: 3,000<br>Each of the other supported Regions: 1,500 | Yes | The maximum number of SendTaskSuccess calls that you can make per second. |
| SendTaskSuccess throttle token refill rate per second | us-east-1: 500<br>us-west-2: 500<br>eu-west-1: 500<br>Each of the other supported Regions: 300 | Yes | The token refill rate per second of SendTaskSuccess calls. |
| StartExecution throttle token bucket size | us-east-1: 1,300<br>us-west-2: 1,300<br>eu-west-1: 1,300<br>Each of the other supported Regions: 800 | Yes | The maximum number of StartExecution calls that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StartExecution throttle token refill rate per second | us-east-1: 300<br>us-west-2: 300<br>eu-west-1: 300<br>Each of the other supported Regions: 150 | Yes | The token refill rate per second of StartExecution calls. |
| StartExpressExecution throttle token bucket size | Each supported Region: 6,000 | Yes | The maximum number of StartExpressExecution calls that you can make per second. |
| StartExpressExecution throttle token refill rate per second | Each supported Region: 6,000 | Yes | The token refill rate per second of StartExpressExecution calls. |
| StartMapRun throttle token bucket size | Each supported Region: 25 | No | The maximum number of StartMapRun calls that you can make per second. |
| StartMapRun throttle token refill rate per second | Each supported Region: 25 | No | The token refill rate per second of StartMapRun calls |
| State machine aliases per state machine | Each supported Region: 100 | Yes | The maximum number of aliases that can be created for a state machine. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StateTransition throttle token bucket size | us-east-1: 5,000<br>us-west-2: 5,000<br>eu-west-1: 5,000<br>Each of the other supported Regions: 800 | Yes | The maximum number of StateTransition calls that you can make per second. |
| StateTransition throttle token refill rate per second | us-east-1: 5,000<br>us-west-2: 5,000<br>eu-west-1: 5,000<br>Each of the other supported Regions: 800 | Yes | The token refill rate per second of StateTransition calls. |
| StopExecution throttle token bucket size | us-east-1: 1,000<br>us-west-2: 1,000<br>eu-west-1: 1,000<br>Each of the other supported Regions: 500 | Yes | The maximum number of StopExecution calls that you can make per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| StopExecution throttle token refill rate per second | us-east-1: 200<br>us-west-2: 200<br>eu-west-1: 200<br>Each of the other supported Regions: 25 | Yes | The token refill rate per second of StopExecution calls. |
| TagResource throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of TagResource calls that you can make per second. |
| TagResource throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of TagResource calls. |
| Task execution time | Each supported Region: 1 | No | The maximum time (in years) that a task can take to complete an execution. |
| Task retention time in queue | Each supported Region: 1 | No | The maximum time (in years) that Step Functions keeps a task in the queue. |
| TestState throttle token bucket size | Each supported Region: 50 | No | The maximum number of TestState calls that you can make per second. |
| TestState throttle token refill rate per second | Each supported Region: 10 | No | The token refill rate per second of TestState calls. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| UntagResource throttle token bucket size | Each supported Region: 200 | Yes | The maximum number of UntagResource calls that you can make per second. |
| UntagResource throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of UntagResource calls. |
| UpdateMapRun throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of UpdateMapRun calls that you can make per second. |
| UpdateMapRun throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for UpdateMapRun calls. |
| UpdateStateMachine throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of UpdateStateMachine calls that you can make per second. |
| UpdateStateMachine throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second of UpdateStateMachine calls. |
| UpdateStateMachineAlias throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of UpdateStateMachineAlias calls that you can make per second. |
| UpdateStateMachineAlias throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for UpdateStateMachineAlias calls. |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| ValidateStateMachineDefinition throttle token bucket size | Each supported Region: 100 | Yes | The maximum number of ValidateStateMachineDefinition calls that you can make per second. |
| ValidateStateMachineDefinition throttle token refill rate per second | Each supported Region: 1 | Yes | The token refill rate per second for ValidateStateMachineDefinition calls. |

For more information, see Quotas in the *AWS Step Functions Developer Guide*.

# AWS Storage Gateway endpoints and quotas

The following are the service endpoints and service quotas for this service. To connect to an AWS service, you use an endpoint. In addition to the standard AWS endpoints, some AWS services offer FIPS endpoints in selected Regions. For more information, see AWS service endpoints. Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

## Service endpoints

### Storage Gateway

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | storagegateway.us-east-2.amazonaws.com | HTTPS | |
| | | storagegateway-fips.us-east-2.api.aws | HTTPS | |
| | | storagegateway.us-east-2.api.aws | HTTPS | |
| | | | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | storagegateway-fips.us-east-2.amazonaws.com | | |
| US East (N. Virginia) | us-east-1 | storagegateway.us-east-1.amazonaws.com<br>storagegateway-fips.us-east-1.api.aws<br>storagegateway-fips.us-east-1.amazonaws.com<br>storagegateway.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (N. California) | us-west-1 | storagegateway.us-west-1.amazonaws.com<br>storagegateway.us-west-1.api.aws<br>storagegateway-fips.us-west-1.amazonaws.com<br>storagegateway-fips.us-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | storagegateway.us-west-2.amazonaws.com<br>storagegateway-fips.us-west-2.amazonaws.com<br>storagegateway.us-west-2.api.aws<br>storagegateway-fips.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | storagegateway.af-south-1.amazonaws.com<br>storagegateway.af-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | storagegateway.ap-east-1.amazonaws.com<br>storagegateway.ap-east-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | storagegateway.ap-south-2.amazonaws.com<br>storagegateway.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | storagegateway.ap-southeast-3.amazonaws.com<br>storagegateway.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | storagegateway.ap-southeast-5.amazonaws.com<br>storagegateway.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | storagegateway.ap-southeast-4.amazonaws.com<br>storagegateway.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | storagegateway.ap-south-1.amazonaws.com<br>storagegateway.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | storagegateway.ap-southeast-6.amazonaws.com<br>storagegateway.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | storagegateway.ap-northeast-3.amazonaws.com<br>storagegateway.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | storagegateway.ap-northeast-2.amazonaws.com<br>storagegateway.ap-northeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | storagegateway.ap-southeast-1.amazonaws.com<br>storagegateway.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | storagegateway.ap-southeast-2.amazonaws.com<br>storagegateway.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | storagegateway.ap-east-2.amazonaws.com<br>storagegateway.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | storagegateway.ap-southeast-7.amazonaws.com<br>storagegateway.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | storagegateway.ap-northeast-1.amazonaws.com<br>storagegateway.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | storagegateway.ca-central-1.amazonaws.com<br>storagegateway.ca-central-1.api.aws<br>storagegateway-fips.ca-central-1.amazonaws.com<br>storagegateway-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | storagegateway.ca-west-1.amazonaws.com | HTTPS | |
| | | storagegateway-fips.ca-west-1.api.aws | HTTPS | |
| | | storagegateway-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | storagegateway.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | storagegateway.eu-central-1.amazonaws.com | HTTPS | |
| | | storagegateway.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | storagegateway.eu-west-1.amazonaws.com | HTTPS | |
| | | storagegateway.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | storagegateway.eu-west-2.amazonaws.com | HTTPS | |
| | | storagegateway.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | storagegateway.eu-south-1.amazonaws.com | HTTPS | |
| | | storagegateway.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | storagegateway.eu-west-3.amazonaws.com | HTTPS | |
| | | storagegateway.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | storagegateway.eu-south-2.amazonaws.com | HTTPS | |
| | | storagegateway.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | storagegateway.eu-north-1.amazonaws.com | HTTPS | |
| | | storagegateway.eu-north-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | storagegateway.eu-central-2.amazonaws.com | HTTPS |
| | | storagegateway.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | storagegateway.il-central-1.amazonaws.com | HTTPS |
| | | storagegateway.il-central-1.api.aws | HTTPS |
| Mexico (Central) | mx-central-1 | storagegateway.mx-central-1.amazonaws.com | HTTPS |
| | | storagegateway.mx-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | storagegateway.me-south-1.amazonaws.com | HTTPS |
| | | storagegateway.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | storagegateway.me-central-1.amazonaws.com | HTTPS |
| | | storagegateway.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | storagegateway.sa-east-1.amazonaws.com | HTTPS |
| | | storagegateway.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | storagegateway.us-gov-east-1.amazonaws.com | HTTPS |
| | | storagegateway-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | storagegateway-fips.us-gov-east-1.api.aws | HTTPS |
| | | storagegateway.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | storagegateway.us-gov-west-1.amazonaws.com<br>storagegateway-fips.us-gov-west-1.amazonaws.com<br>storagegateway.us-gov-west-1.api.aws<br>storagegateway-fips.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Storage Gateway hardware appliance regions

The Storage Gateway hardware appliance is available for shipping worldwide where it is legally allowed and permitted for exporting by the US government.

Storage Gateway hardware appliance is supported in the following AWS Regions.

- US East (Ohio)
- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Mumbai)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Canada (Central)
- Europe (Frankfurt)
- Europe (Ireland)
- Europe (London)
- Europe (Paris)

- Europe (Stockholm)
- South America (São Paulo)

## Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Cached volume gateway Cache Maximum in TiB | Each supported Region: 64 | No | Maximum cache size for Cached Volume Gateway |
| Cached volume gateway Cache Minimum in GiB | Each supported Region: 150 | No | Minimum cache size for Cached Volume Gateway |
| Cached volume gateway Upload Buffer Maximum in TiB | Each supported Region: 2 | No | Maximum upload buffer size for Cached Volume Gateway |
| Cached volume gateway Upload Buffer Minimum in GiB | Each supported Region: 150 | No | Minimum upload buffer size for Cached Volume Gateway |
| Cached volume size in TiB | Each supported Region: 32 | No | Maximum size of a cached volume |
| Cached volumes per gateway | Each supported Region: 32 | No | Maximum number of cached volumes per gateway |
| File gateway Cache Maximum in TiB | Each supported Region: 64 | No | Maximum cache size for File Gateway |
| File gateway Cache Minimum in GiB | Each supported Region: 150 | No | Minimum cache size for File Gateway |
| File shares per S3 bucket | Each supported Region: 1 | No | Maximum number of file shares per Amazon S3 bucket |

| Name | Default | Adjuse | Description |
|---|---|---|---|
| File shares per gateway | Each supported Region: 50 | No | Maximum number of file shares per gateway |
| File size | Each supported Region: 5 Terabytes | No | The maximum size of an individual file, which is the maximum size of an individual object in Amazon S3 |
| Max size of a virtual tape in TiB | Each supported Region: 15 | No | Maximum size of a virtual tape |
| Max virtual tapes in a VTL | Each supported Region: 1,500 | No | Maximum number of virtual tapes for a virtual tape library (VTL) |
| Minimum size of a virtual tape in GiB | Each supported Region: 100 | No | Minimum size of a virtual tape |
| Path length | Each supported Region: 1,024 Bytes | No | Maximum path length |
| Size of all cached volumes per gateway in TiB | Each supported Region: 1,024 | No | Total size of all cached volumes for a gateway |
| Size of all stored volumes per gateway in TiB | Each supported Region: 512 | No | Total size of all stored volumes for a gateway |
| Stored volume gateway Upload Buffer Maximum in TiB | Each supported Region: 2 | No | Maximum upload buffer size for Stored Volume Gateway |
| Stored volume gateway Upload Buffer Minimum in GiB | Each supported Region: 150 | No | Minimum upload buffer size for Stored Volume Gateway |

| Name | Default | Adjuste | Description |
|---|---|---|---|
| Stored volume size in TiB | Each supported Region: 16 | No | Maximum size of a stored volume |
| Stored volumes per gateway | Each supported Region: 32 | No | Maximum number of stored volumes per gateway |
| Tape gateway Cache Maximum in TiB | Each supported Region: 64 | No | Maximum cache size for Tape Gateway |
| Tape gateway Cache Minimum in GiB | Each supported Region: 150 | No | Minimum cache size for Tape Gateway |
| Tape gateway Upload Buffer Maximum in TiB | Each supported Region: 2 | No | Maximum upload buffer size for Tape Gateway |
| Tape gateway Upload Buffer Minimum in GiB | Each supported Region: 150 | No | Minimum upload buffer size for Tape Gateway |
| Total size of tapes in a virtual tape library in PiB | Each supported Region: 1 | No | Total size of all tapes in a virtual tape library (VTL) |

For more information, see Storage Gateway quotas in the *AWS Storage Gateway User Guide*.

## AWS Support endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | support.us-west-2.amazonaws.com | HTTPS | |
| | | support-fips.us-west-2.api.aws | HTTPS | |
| | | support-fips.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Osaka) | ap-northeast-3 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | support.eu-west-1.amazonaws.com | HTTPS | |
| | | support-fips.eu-west-1.api.aws | HTTPS | |
| | | support-fips.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | support.us-east-1.amazonaws.com | HTTPS | |
| | | support-fips.us-east-1.api.aws | HTTPS | |
| | | support-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | support.us-east-1.amazonaws.com | HTTPS |
| | | support-fips.us-east-1.api.aws | HTTPS |
| | | support-fips.us-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | support.us-gov-west-1.amazonaws.com | HTTPS |
| | | support-fips.us-gov-west-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | support.us-gov-west-1.amazonaws.com | HTTPS |
| | | support-fips.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjuste | Description |
|---|---|---|---|
| AWS Support API operations | Each supported Region: 5 | No | The maximum number of AWS Support API operations that you can perform per second. |
| AWS Trusted Advisor API operations | Each supported Region: 100 | No | The maximum number of AWS Trusted Advisor API operations that you can perform per second. |
| Number of AWS Support cases that you can create | Each supported Region: 10 | No | The maximum number of AWS Support cases that you can create per hour. |

For more information, see the AWS Support User Guide.

## AWS Support App in Slack endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | supportapp.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | supportapp.us-west-2.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | supportapp.eu-west-1.amazonaws.com | HTTPS | |

# AWS Systems Manager endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Related services**

For information about endpoints and quotas for related services, see the following topics:

- AWS AppConfig endpoints and quotas
- AWS Systems Manager Incident Manager endpoints and quotas
- Systems Manager Quick Setup endpoints and quotas
- AWS Systems Manager for SAP endpoints and quotas

## Service endpoints for Systems Manager

**Note**

In addition to the `ssm.*` endpoints documented in the following table, your managed nodes must also allow HTTPS (port 443) outbound traffic to the following endpoints.

- `ec2messages.*`
- `ssmmessages.*`

For more information, see Reference: ec2messages, ssmmessages, and other API operations in the *AWS Systems Manager User Guide*.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ssm.us-east-2.amazonaws.com | HTTPS |
| | | ssm-fips.us-east-2.amazonaws.com | HTTPS |
| | | ssm-fips.us-east-2.api.aws | HTTPS |
| | | ssm.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | ssm.us-east-1.amazonaws.com | HTTPS |
| | | ssm-fips.us-east-1.amazonaws.com | HTTPS |
| | | ssm-fips.us-east-1.api.aws | HTTPS |
| | | ssm.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ssm.us-west-1.amazonaws.com | HTTPS |
| | | ssm-fips.us-west-1.amazonaws.com | HTTPS |
| | | ssm-fips.us-west-1.api.aws | HTTPS |
| | | ssm.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ssm.us-west-2.amazonaws.com | HTTPS |
| | | ssm-fips.us-west-2.amazonaws.com | HTTPS |
| | | ssm-fips.us-west-2.api.aws | HTTPS |
| | | ssm.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | ssm.af-south-1.amazonaws.com | HTTPS |
| | | ssm.af-south-1.api.aws | HTTPS |
| Asia Pacific | ap-east-1 | ssm.ap-east-1.amazonaws.com | HTTPS |
| | | ssm.ap-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | ssm.ap-south-2.amazonaws.com<br>ssm.ap-south-2.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | ssm.ap-southeast-3.amazonaws.com<br>ssm.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | ssm.ap-southeast-5.amazonaws.com<br>ssm.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | ssm.ap-southeast-4.amazonaws.com<br>ssm.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | ssm.ap-south-1.amazonaws.com<br>ssm.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | ssm.ap-southeast-6.amazonaws.com<br>ssm.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | ssm.ap-northeast-3.amazonaws.com<br>ssm.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | ssm.ap-northeast-2.amazonaws.com | HTTPS |
| | | ssm.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | ssm.ap-southeast-1.amazonaws.com | HTTPS |
| | | ssm.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ssm.ap-southeast-2.amazonaws.com | HTTPS |
| | | ssm.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | ssm.ap-east-2.amazonaws.com | HTTPS |
| | | ssm.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | ssm.ap-southeast-7.amazonaws.com | HTTPS |
| | | ssm.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ssm.ap-northeast-1.amazonaws.com | HTTPS |
| | | ssm.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | ssm.ca-central-1.amazonaws.com | HTTPS |
| | | ssm-fips.ca-central-1.amazonaws.com | HTTPS |
| | | ssm-fips.ca-central-1.api.aws | HTTPS |
| | | ssm.ca-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | ssm.ca-west-1.amazonaws.com | HTTPS | |
| | | ssm-fips.ca-west-1.amazonaws.com | HTTPS | |
| | | ssm-fips.ca-west-1.api.aws | HTTPS | |
| | | ssm.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | ssm.eu-central-1.amazonaws.com | HTTPS | |
| | | ssm.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | ssm.eu-west-1.amazonaws.com | HTTPS | |
| | | ssm.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | ssm.eu-west-2.amazonaws.com | HTTPS | |
| | | ssm.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | ssm.eu-south-1.amazonaws.com | HTTPS | |
| | | ssm.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | ssm.eu-west-3.amazonaws.com | HTTPS | |
| | | ssm.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | ssm.eu-south-2.amazonaws.com | HTTPS | |
| | | ssm.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | ssm.eu-north-1.amazonaws.com | HTTPS | |
| | | ssm.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | ssm.eu-central-2.amazonaws.com | HTTPS | |
| | | ssm.eu-central-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Israel (Tel Aviv) | il-central-1 | ssm.il-central-1.amazonaws.com<br>ssm.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | ssm.mx-central-1.amazonaws.com<br>ssm.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | ssm.me-south-1.amazonaws.com<br>ssm.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | ssm.me-central-1.amazonaws.com<br>ssm.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | ssm.sa-east-1.amazonaws.com<br>ssm.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | ssm.us-gov-east-1.amazonaws.com<br>ssm.us-gov-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | ssm.us-gov-west-1.amazonaws.com<br>ssm.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

## Service quotas

The following sections list and describe the quotas for Systems Manager, grouped by tool or feature area.

Unless otherwise specified, each quota applies to an individual AWS Region in each AWS account. For example, the default quota for the number of State Manager associations is 2,000. This means that in the AWS account 123456789012, you can create 2,000 associations in the US East (N. Virginia) Region, 2,000 in the US East (Ohio) Region, and so on.

**Topics**


- Service quotas for Application Manager
- Service quotas for Automation
- Service quotas for Distributor
- Service quotas for Documents
- Service quotas for Explorer
- Service quotas for Fleet Manager / AWS Systems Manager GUI Connect
- Service quotas for Inventory
- Service quotas for Maintenance Windows
- Service quotas for Managed nodes
- Service quotas for OpsCenter
- Service quotas for Parameter Store
- Service quotas for Patch Manager
- Service quotas for Session Manager
- Service quotas for State Manager


## Service quotas for Application Manager

| Resource | Default |
| --- | --- |
| Maximum number of applications in Application Manager | 100<br><br>When you add an application in Application Manager, Systems Manager automatically creates a resource group to organize all of the resources for that application. The maximum number of applications is |

| Resource | Default |
| --- | --- |
| | based on the underlying quota for AWS Resource Groups. |
| Maximum number of AWS resources assigned to an application | For applications based on AWS CloudFormation stacks: 200<br><br>For applications based on AWS Resource Groups: Unlimited |

## Service quotas for Automation

| Resource | Default |
| --- | --- |
| Concurrently running automations | 100<br><br>This quota can be increased up to 500 by enabling adaptive concurrency.<br><br>Additionally, you can run up to 400 concurrent automations with blocking actions. Blocking actions include `aws:approve`, `aws:pause`, and `aws:sleep`. If you attempt to run more automations than this, Systems Manager adds the additional automations to a queue and displays a status of `Pending`.<br><br>For more information about adaptive concurrency, see Allowing Automation to adapt to your concurrency needs in the *AWS Systems Manager User Guide*. |
| Automation queue | 5,000 |

| Resource | Default |
| --- | --- |
| | If you attempt to run more automatio ns than the concurrent automation limit, subsequent automations are added to the Automation queue. When an automation completes (or reaches a terminal state), the first automation in the queue is started. |
| Concurrently running rate control automations | 25<br><br>If you attempt to run more rate control automations than the concu rrent rate control automation limit, Systems Manager adds the subse quent rate control automations to the queue and displays a status of `Pending`. |
| Rate control automation queue | 1,000<br><br>If you attempt to run more automatio ns than the concurrent rate control automation limit, subsequent automations are added to the queue. When an automation completes (or reaches a terminal state), the first automation in the queue is started. |

| Resource | Default |
| --- | --- |
| Number of levels of nested automation | 5<br><br>A parent-level Automation runbook can start a child-level Automation runbook. This represents one level of nested automation. The child-level Automation runbook can start another Automation runbook, resulting in two levels of nested au tomation. This can continue up to a maximum of five levels below the top-level parent Automation runbook. |
| Number of days an automation execution history is stored in the system | 30 |
| Number of days an automation variable is stored in the system | 30 |
| Additional automation executions that can be queued | 1,000 |
| Maximum duration an automation execution can run in the context of a user | 12 hours<br><br>If you expect an automation to run longer than 12 hours, then you must run the automation by using a service role (or assume role). |
| Maximum `executeScript` action run time | 10 minutes<br><br>Each `executeScript` action can run up to a maximum duration of 10 minutes. |

| Resource | Default |
| --- | --- |
| Maximum `executeScript` action maximum output | 100 KB |
| Maximum `invokeLambdaFunction` action run time | 5 minutes<br><br>Each `invokeLambdaFunction` action can run up to a maximum duration of five (5) minutes. |
| Maximum `invokeLambdaFunction` action output | 200 KB |
| Number of Automation runbook attachments | 5 per runbook |
| Size of an Automation runbook attachment | 256 MB per attachment, per runbook |
| Transactions per second for the StartAutomationExecution API action | 1 |
| Transactions per second for the DescribeAutomationStepExecutions API action | 3 |
| Transactions per second for the GetAutomationExecution API action | 3 |
| Transactions per second for the DescribeAutomationExecutions API action | 3 |
| Transactions per second for the SendAutomationSignal API action | 3 |
| Transactions per second for the StopAutomationExecution API action | 1 |
| Transactions per second for the StartExecutionPreview API action | 1 |

| Resource | Default |
| --- | --- |
| Transactions per second for the GetExecutionPreview API action | 1 |

## Service quotas for Distributor

| Resource | Default |
| --- | --- |
| Maximum number of attachments in a Distributor package | 20 |
| Maximum size per attachment in a Distributor package | 1 GB |
| Maximum number of files in a Distributor package | 1,000 |
| Maximum number of Distributor packages | 500 |
| Maximum number of package versions per Distributor package | 25 |
| Maximum package size in Distributor | 20 GB |
| Maximum package manifest size in Distributor | 64 KB |

## Service quotas for Documents

| Resource | Default |
| --- | --- |
| Document size | 64 KB per document |
| Total documents | |

| Resource | Default |
| --- | --- |
| | 500 |
| Document versions | 1,000 per document |
| Privately shared Systems Manager document | A single SSM document can be shared with a maximum of 1,000 AWS accounts. |
| Publicly shared Systems Manager document | 5<br><br>Each AWS account can publicly share a maximum of five documents. |
| Maximum number of favorites per document type | 20 |

## Service quotas for Explorer

| Resource | Default |
| --- | --- |
| Maximum number of resource data syncs | 5 |

## Service quotas for Fleet Manager / AWS Systems Manager GUI Connect

| Resource | Default |
| --- | --- |
| Maximum number of Systems Manager GUI Connect concurrent Remote Desktop sessions | 5 [1]<br><br>Service quota increase requests up to 50 are automatically approved. Service quota increases can take up to two and a half hours to take effect. |

| Resource | Default |
| --- | --- |
| Maximum duration of an Systems Manager GUI Connect Remote Desktop session | 60 minutes |
| Maximum idle time for a Systems Manager GUI Connect Remote Desktop session before connection timeout | 10 minutes |
| Maximum connection persistence for a Systems Manager GUI Connect Remote Desktop session | 60 minutes (maximum connection duration) or 10 minutes (idle timeout limit).<br><br>A connection persists until the maximum connection duration (60 minutes) or idle timeout limit (10 minutes) is met. The connection persists after IAM credentials expire if the connection duration limits are not met. |

[1] The standard license for Windows Server allows for two concurrent RDP connections. To support more connections, you must purchase additional Client Access Licenses (CALs) from Microsoft or Microsoft Remote Desktop Services licenses from AWS. For more information on supplemental licensing, see the following topics:

- Client Access Licenses and Management Licenses on the Microsoft website
- Use License Manager user-based subscriptions for supported software products in the *License Manager User Guide*

## Service quotas for Inventory

| Resource | Default |
| --- | --- |
| Maximum size of an inventory data item that can be sent in a single `PutInventory` API request, per managed node | 1024 KB |
| Maximum number of resource data syncs | 5 |
| Maximum size of inventory data collected per managed node, per call | 1 MB<br><br>This maximum adequately supports most inventory collection scenarios. When this quota is reached, no new inventory data is collected for the node. Inventory data previously collected is stored until the expiration. |
| Maximum size of inventory data collected per node, per day | 5,000 KB<br><br>When this quota is reached, no new inventory data is collected for the instance. Inventory data previously collected is stored until the expiration. |
| Number of custom inventory types | 20 |
| Maximum size of a custom inventory type | 200 KB<br><br>This is the maximum size of the type, not the inventory collected. |

| Resource | Default |
| --- | --- |
| Maximum number of attributes in a custom inventory type | 50 |

| Resource | Default |
| --- | --- |
| Length of inventory data retention | 30 days<br><br>AWS Systems Manager Inventory retains managed node data for 30 days after the most recent inventory update. This retention quota applies to all managed node states, including stopped instances. The retention rules work this way:<br><br>• If a node remains in a stopped state or loses connection for 30 consecutive days without any inventory updates, the inventory data is automatically deleted after the 30-day period.<br>• If that node becomes active again before the 30-day period expires (for example, if a stopped node is started) and new inventory data is reported successfully, the 30-day retention period resets from that point.<br><br>If you need to store inventory data longer than 30 days, you can use AWS Config to record history or periodically query and upload the data to an Amazon S3 bucket. For more information, see Recording Software Configuration for Managed Instances in the *AWS Config Developer Guide*. |

## Service quotas for Maintenance Windows

| Resource | Default |
| --- | --- |
| Maintenance windows | 50 |
| Tasks per maintenance window | 20 |
| Targets per maintenance window | 100 |
| Targets per task | 10 |
| Concurrent executions of maintenance windows | 5 |

## Service quotas for Managed nodes

| Resource | Default |
| --- | --- |
| Maximum number of managed nodes (Amazon EC2 and hybrid) in a fleet | 2,400<br><br>Approximate maximum number of nodes managed by Systems Manager (per AWS account per Region). Utilization up to the applied account level quota is considered safe. Exceeding this quota may cause ins tances to stop communicating with Systems Manager, but actual capacity varies.<br><br>**AWS default quota value**: The default fleet size quota is set at 2,400 managed nodes in Systems Manager. This is the default maximum safe |

| Resource | Default |
| --- | --- |
| | limit for an AWS account. We do not recommend scaling past this without a limit increase because instances could stop communicating with Systems Manager. If you pass this limit, instances might be able to connect to Systems Manager, especially if you stagger node launches, but it's not certain.<br><br>**Applied account-level quota value**: The applied quota is the current regional limit for this account. This may be the default quota or an increased limit that has been approved by Support upon customer request. Exceeding the limit may cause instances to stop communicating with Systems Manager because connection attempts could exceed approved throughput.<br><br>**Utilization**: The utilized quota represents the approximate number of managed nodes currently active and communicating with Systems Manager. It is calculated based on the number of UpdateInstanceInformation API calls made within a five minute period of time. This number changes when instances are launched, terminated, or experience connectivity issues. The number also changes when SSM Agent is restarted. Customers may see greater |

| Resource | Default |
| --- | --- |
| | than 100% utilization if instances are configured to call `UpdateInstanceInformation` API more frequently than the default five-minute interval or if instances launched within a five-minute window are staggered, resulting in distributed calls to the UpdateInstanceInformation API. |
| Maximum number of hybrid-activated machines in a hybrid and multicloud environment | Standard instances: 1,000<br><br>Advanced instances: Advanced instances are available on a pay-per-use basis. Advanced instances also enable you to connect to your non-EC2 machines by using AWS Systems Manager Session Manager.<br><br>For more information about activating non-EC2 machines for use in your hybrid and multicloud environment, see Setting up Systems Manager for hybrid and multicloud environments in the *AWS Systems Manager User Guide*. For more information about enabling advanced instances, see Configuring instance tiers. |

## Service quotas for OpsCenter

| Resource | Default |
| --- | --- |
| | 500,000 |

| Resource | Default |
| --- | --- |
| Maximum number of OpsItems (including Open *and* Resolved OpsItems) | **Important**<br>OpsItems that are created by an integration with AWS Security Hub CSPM are *not* currently limited by this maximum quota. It is therefore possible for Security Hub CSPM alerts to create more than 500,000 chargeable OpsItems in an account. For high-production environments, we therefore recommend limiting the scope of Security Hub CSPM findings to high severity issues only.<br>For more information about OpsCenter integration with Security Hub CSPM and OpsItems pricing in AWS Systems Manager, see Understanding OpsCenter integration with AWS Security Hub CSPM in the *AWS Systems Manager User Guide*. |
| Maximum number of OpsItems per AWS account per month | 10,000 |
| Maximum operational data value size | 20 KB |

| Resource | Default |
| --- | --- |
| Maximum number of associated Automation runbooks per OpsItem | 10 |
| Maximum number of Automation runbook execution s stored in operational data under a single associated runbook | 10 |
| Maximum number of related resources you can specify per OpsItem | 100 |
| Maximum number of related OpsItems you can specify per OpsItem | 10 |
| Maximum length of a deduplication string | 512 characters |

## Service quotas for Parameter Store

| Resource | Default |
| --- | --- |
| Maximum number of parameters | Standard parameters: 10,000<br>Advanced parameters: 100,000 [1] |
| Maximum size for parameter value | Standard parameter: 4 KB<br>Advanced parameter: 8 KB [1] |
| Maximum number of parameter policies per advanced parameter | 10 [1] |

| Resource | Default |
| --- | --- |
| Maximum number of parameter versions retained | 100 |
| Transactions per second for the following API Actions:<br>• GetParameter<br>• GetParameters<br>• GetParametersByPath | • Default: 40<br>This default maximum of 40 TPS is shared by all three API actions.<br>• Higher throughput enabled: [2, 3]<br>• `GetParameter` : 10,000<br>• `GetParameters` : 1,000<br>• `GetParametersByPath` : 100 |
| Transactions per second for the following API actions:<br>• DeleteParameter<br>• DeleteParameters | • Default: 3<br>• Higher throughput enabled: 5 [2] |
| Transactions per second for the following API actions:<br>• DescribeParameters<br>• GetParameterHistory<br>• LabelParameterVersion<br>• UnlabelParameterVersion<br>• PutParameter | • Default: 3<br>• Higher throughput enabled: 10 [2] |

[1] Enabling the advanced parameter tier incurs a charge on your AWS account. For more information, see Managing parameter tiers in the *AWS Systems Manager User Guide*.

[2] You can raise the maximum transactions per second (TPS) for this API action to support applications and workloads that need concurrent access to multiple parameters. Increasing the TPS quota incurs a charge on your AWS account. For more information, see Increasing or resetting Parameter Store throughput in the *AWS Systems Manager User Guide*.

[3] Throughput for `SecureString` parameters might be further limited by AWS Key Management Service (AWS KMS) throughput limits depending on the Region. For more information about AWS KMS limits, see Request quotas in the *AWS Key Management Service Developer Guide*.

## Service quotas for Patch Manager

| Resource | Default |
|---|---|
| Patch baselines | 50 |
| Patch groups per patch baseline | 25 |
| Operation history retention | Most recent 150 operations |

## Service quotas for Session Manager

| Resource | Default |
|---|---|
| Transactions per second for the DescribeSessions API action | 6 |
| Transactions per second for the GetConnectionStatus API action | 50 |
| Transactions per second for the ResumeSession API action | 6 |
| Transactions per second for the StartSession API action | 3 |

| Resource | Default |
| --- | --- |
| Transactions per second for the TerminateSession API action | 6 |
| Idle time before session termination | Default: 20 minutes<br><br>Configurable to between 1 and 60 minutes. |
| Execution history retention | 30 days<br><br>The history of each command is available for up to 30 days. In additi on, you can store a copy of all log files in Amazon Simple Storage Service or have an audit trail of all API calls in AWS CloudTrail. |

## Service quotas for State Manager

| Resource | Default |
| --- | --- |
| Transactions per second (TPS) for the CreateAss ociation API action | 3 |
| Transactions per second (TPS) for the CreateAss ociationBatch API action | 1 |
| Transactions per second (TPS) for the DeleteAss ociation API action | 2 |
| Transactions per second (TPS) for the DescribeA ssociation API action | 2 |
| Transactions per second (TPS) for the DescribeA ssociationExecutions API action | 4 |

| Resource | Default |
| --- | --- |
| Transactions per second (TPS) for the DescribeAssociationExecutionTargets API action | 4 |
| Transactions per second (TPS) for the DescribeEffectiveInstanceAssociations API action | 1 |
| Transactions per second (TPS) for the DescribeInstanceAssociationsStatus API action | 2 |
| Transactions per second (TPS) for the ListAssociations API action | 4 |
| Transactions per second (TPS) for the ListAssociationVersions API action | 4 |
| Transactions per second (TPS) for the StartAssociationsOnce API action | 2 |
| Transactions per second (TPS) for the UpdateAssociation API action | 3 |
| Transactions per second (TPS) for the UpdateAssociationStatus API action | 3 |
| Maximum number of associations | 2,000 |
| Maximum number of versions per association | 1,000 |
| Maximum number of associations targeting a single managed node | 20 |

# AWS Systems Manager for SAP endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | ssm-sap.us-east-2.amazonaws.com | HTTPS |
| | | ssm-sap-fips.us-east-2.api.aws | HTTPS |
| | | ssm-sap-fips.us-east-2.amazonaws.com | HTTPS |
| | | ssm-sap.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | ssm-sap.us-east-1.amazonaws.com | HTTPS |
| | | ssm-sap-fips.us-east-1.api.aws | HTTPS |
| | | ssm-sap-fips.us-east-1.amazonaws.com | HTTPS |
| | | ssm-sap.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | ssm-sap.us-west-1.amazonaws.com | HTTPS |
| | | ssm-sap-fips.us-west-1.api.aws | HTTPS |
| | | ssm-sap-fips.us-west-1.amazonaws.com | HTTPS |
| | | ssm-sap.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | ssm-sap.us-west-2.amazonaws.com | HTTPS |
| | | ssm-sap-fips.us-west-2.api.aws | HTTPS |
| | | ssm-sap-fips.us-west-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | ssm-sap.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | ssm-sap.af-south-1.amazonaws.com | HTTPS |
| | | ssm-sap.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | ssm-sap.ap-east-1.amazonaws.com | HTTPS |
| | | ssm-sap.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | ssm-sap.ap-south-2.amazonaws.com | HTTPS |
| | | ssm-sap.ap-south-2.api.aws | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | ssm-sap.ap-southeast-3.amazonaws.com | HTTPS |
| | | ssm-sap.ap-southeast-3.api.aws | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | ssm-sap.ap-southeast-4.amazonaws.com | HTTPS |
| | | ssm-sap.ap-southeast-4.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ssm-sap.ap-south-1.amazonaws.com | HTTPS |
| | | ssm-sap.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | ssm-sap.ap-northeast-3.amazonaws.com | HTTPS |
| | | ssm-sap.ap-northeast-3.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | ssm-sap.ap-northeast-2.amazonaws.com | HTTPS |
| | | ssm-sap.ap-northeast-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | ssm-sap.ap-southeast-1.amazonaws.com | HTTPS |
| | | ssm-sap.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ssm-sap.ap-southeast-2.amazonaws.com | HTTPS |
| | | ssm-sap.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ssm-sap.ap-northeast-1.amazonaws.com | HTTPS |
| | | ssm-sap.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | ssm-sap.ca-central-1.amazonaws.com | HTTPS |
| | | ssm-sap-fips.ca-central-1.api.aws | HTTPS |
| | | ssm-sap-fips.ca-central-1.amazonaws.com | HTTPS |
| | | ssm-sap.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | ssm-sap.eu-central-1.amazonaws.com | HTTPS |
| | | ssm-sap.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | ssm-sap.eu-west-1.amazonaws.com | HTTPS |
| | | ssm-sap.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | ssm-sap.eu-west-2.amazonaws.com | HTTPS |
| | | ssm-sap.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | ssm-sap.eu-south-1.amazonaws.com | HTTPS |
| | | ssm-sap.eu-south-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Paris) | eu-west-3 | ssm-sap.eu-west-3.amazonaws.com | HTTPS |
| | | ssm-sap.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | ssm-sap.eu-south-2.amazonaws.com | HTTPS |
| | | ssm-sap.eu-south-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | ssm-sap.eu-north-1.amazonaws.com | HTTPS |
| | | ssm-sap.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | ssm-sap.eu-central-2.amazonaws.com | HTTPS |
| | | ssm-sap.eu-central-2.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | ssm-sap.il-central-1.amazonaws.com | HTTPS |
| | | ssm-sap.il-central-1.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | ssm-sap.me-south-1.amazonaws.com | HTTPS |
| | | ssm-sap.me-south-1.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | ssm-sap.me-central-1.amazonaws.com | HTTPS |
| | | ssm-sap.me-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | ssm-sap.sa-east-1.amazonaws.com | HTTPS |
| | | ssm-sap.sa-east-1.api.aws | HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| SAP applications per Region in an AWS account | 10 | Yes | The maximum number of SAP applications that you can register with AWS Systems Manager for SAP per Region in an AWS account. |
| Components per SAP application | 20 | Yes | The maximum number of `ssm-sap` components that you can register per SAP application registered with AWS Systems Manager for SAP. |
| Databases per component | 20 | Yes | The maximum number of `ssm-sap` databases that you can register per `ssm-sap` component. |

# Amazon Textract endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | textract.us-east-2.amazonaws.com | HTTPS | |
| | | textract-fips.us-east-2.api.aws | HTTPS | |
| | | textract-fips.us-east-2.amazonaws.com | HTTPS | |
| | | textract.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | textract.us-east-1.amazonaws.com | HTTPS | |
| | | textract-fips.us-east-1.api.aws | HTTPS | |
| | | textract-fips.us-east-1.amazonaws.com | HTTPS | |
| | | textract.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | textract.us-west-1.amazonaws.com | HTTPS | |
| | | textract-fips.us-west-1.api.aws | HTTPS | |
| | | textract-fips.us-west-1.amazonaws.com | HTTPS | |
| | | textract.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | textract.us-west-2.amazonaws.com | HTTPS | |
| | | textract-fips.us-west-2.api.aws | HTTPS | |
| | | textract-fips.us-west-2.amazonaws.com | HTTPS | |
| | | textract.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | textract.ap-south-1.amazonaws.com | HTTPS | |
| | | textract.ap-south-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | textract.ap-northeast-2.amazonaws.com | HTTPS |
| | | textract.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | textract.ap-southeast-1.amazonaws.com | HTTPS |
| | | textract.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | textract.ap-southeast-2.amazonaws.com | HTTPS |
| | | textract.ap-southeast-2.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | textract.ca-central-1.amazonaws.com | HTTPS |
| | | textract-fips.ca-central-1.api.aws | HTTPS |
| | | textract-fips.ca-central-1.amazonaws.com | HTTPS |
| | | textract.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | textract.eu-central-1.amazonaws.com | HTTPS |
| | | textract.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | textract.eu-west-1.amazonaws.com | HTTPS |
| | | textract.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | textract.eu-west-2.amazonaws.com | HTTPS |
| | | textract.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | textract.eu-west-3.amazonaws.com | HTTPS |
| | | textract.eu-west-3.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Spain) | eu-south-2 | textract.eu-south-2.amazonaws.com | HTTPS |
| | | textract.eu-south-2.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | textract.us-gov-east-1.amazonaws.com | HTTPS |
| | | textract-fips.us-gov-east-1.api.aws | HTTPS |
| | | textract-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | textract.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | textract.us-gov-west-1.amazonaws.com | HTTPS |
| | | textract-fips.us-gov-west-1.api.aws | HTTPS |
| | | textract-fips.us-gov-west-1.amazonaws.com | HTTPS |
| | | textract.us-gov-west-1.api.aws | HTTPS |

## Service quotas

| Resources | | Regions | | | | | |
|---|---|---|---|---|---|---|---|
| Synchronous Operations | API | US East (N. Virginia) | US West (Oregon) | US East (Ohio) | Europe (Ireland) | Asia Pacific (Mumbai) | Other Regions |
| Transactions per second per account for synchrono | AnalyzeDocument | 10 | 10 | 10 | 5 | 5 | 1 |
| | DetectDocumentText | 25 | 25 | 10 | 5 | 5 | 1 |

| Resources | | Regions | | | | | |
|---|---|---|---|---|---|---|---|
| us operations | AnalyzeExpense | 5 | 5 | 1 | 1 | 1 | 1 |
| | AnalyzeID | 5 | 5 | 1 | 1 | 1 | 1 |
| Asynchronous Operations | API | US East (N. Virginia) | US West (Oregon) | US East (Ohio) | Europe (Ireland) | Asia Pacific (Mumbai) | Other Regions |
| Transactions per second per account for all start (asynchronous) operations | StartDocumentAnalysis | 10 | 10 | 10 | 5 | 5 | 2 |
| | StartDocumentTextDetection | 15 | 15 | 5 | 5 | 5 | 1 |
| | StartExpenseAnalysis | 5 | 5 | 1 | 1 | 1 | 1 |
| | StartLendingAnalysis | 5 | 5 | 1 | 1 | 1 | 1 |

| Resources | | Regions | | | | | |
|---|---|---|---|---|---|---|---|
| Transactions per second per account for all get (asynchronous) operations | GetDocumentAnalysis | 10 | 10 | 10 | 5 | 5 | 5 |
| | GetDocumentTextDetection | 25 | 25 | 10 | 5 | 5 | 5 |
| | GetExpenseAnalysis | 5 | 5 | 5 | 5 | 5 | 5 |
| | GetLendingAnalysis | 25 | 25 | 5 | 5 | 5 | 5 |
| | GetLendingAnalysisSummary | 5 | 5 | 1 | 1 | 1 | 1 |
| Maximum number of asynchronous jobs per account that can simultaneously exist | | 600 | 600 | 100 | 100 | 100 | 100 |

For more information, see Amazon Textract Quotas in the *Amazon Textract Developer Guide*.

## Adapters quotas

- Maximum number of adapters - Total number of adapters allowed are 10. You can have a several adapter versions under a single adapter.

- Maximum AdapterVersions created per month - Number of successful adapter versions that can be created per AWS account per month is 10 which will be reset at the start of every month. Use the Service Quotas console to raise a service quota increase request.
- Maximum in-progress AdapterVersions (analogous to adapter training) per account - 3

For more information, see Amazon Textract Quotas in the *Amazon Textract Developer Guide*.

# Amazon Timestream endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

### Timestream for Live Analytics endpoints

Use the following endpoints to acquire the endpoints for the write API.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | ingest.timestream.us-east-2.amazonaws.com | HTTPS | |
| | | timestream-ingest-fips.us-east-2.api.aws | HTTPS | |
| | | ingest.timestream-fips.us-east-2.amazonaws.com | HTTPS | |
| | | timestream-ingest.us-east-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | ingest.timestream.us-east-1.amazonaws.com<br>timestream-ingest-fips.us-east-1.api.aws<br>ingest.timestream-fips.us-east-1.amazonaws.com<br>timestream-ingest.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | ingest.timestream.us-west-2.amazonaws.com<br>timestream-ingest-fips.us-west-2.api.aws<br>ingest.timestream-fips.us-west-2.amazonaws.com<br>timestream-ingest.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | ingest.timestream.ap-south-1.amazonaws.com<br>timestream-ingest.ap-south-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | ingest.timestream.ap-southeast-2.amazonaws.com<br>timestream-ingest.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | ingest.timestream.ap-northeast-1.amazonaws.com<br>timestream-ingest.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | ingest.timestream.eu-central-1.amazonaws.com<br>timestream-ingest.eu-central-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Ireland) | eu-west-1 | ingest.timestream.eu-west-1.amazonaws.com<br>timestream-ingest.eu-west-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | ingest.timestream.us-gov-west-1.amazonaws.com<br>timestream-ingest.us-gov-west-1.api.aws | HTTPS<br>HTTPS |

Use the following endpoints to acquire the endpoints for the query API.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | query.timestream.us-east-2.amazonaws.com<br>timestream-query-fips.us-east-2.api.aws<br>query.timestream-fips.us-east-2.amazonaws.com<br>timestream-query.us-east-2.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US East (N. Virginia) | us-east-1 | query.timestream.us-east-1.amazonaws.com<br>timestream-query-fips.us-east-1.api.aws<br>query.timestream-fips.us-east-1.amazonaws.com<br>timestream-query.us-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| US West (Oregon) | us-west-2 | query.timestream.us-west-2.amazonaws.com<br>timestream-query-fips.us-west-2.api.aws | HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | query.timestream-fips.us-west-2.amazonaws.com<br><br>timestream-query.us-west-2.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | query.timestream.ap-south-1.amazonaws.com<br><br>timestream-query.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | query.timestream.ap-southeast-2.amazonaws.com<br><br>timestream-query.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | query.timestream.ap-northeast-1.amazonaws.com<br><br>timestream-query.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | query.timestream.eu-central-1.amazonaws.com<br><br>timestream-query.eu-central-1.api.aws | HTTPS<br><br>HTTPS | |
| Europe (Ireland) | eu-west-1 | query.timestream.eu-west-1.amazonaws.com<br><br>timestream-query.eu-west-1.api.aws | HTTPS<br><br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | query.timestream.us-gov-west-1.amazonaws.com<br><br>timestream-query.us-gov-west-1.api.aws | HTTPS<br><br>HTTPS | |

For more information, see Using the API in the *Amazon Timestream Developer Guide*.

## Timestream for InfluxDB endpoints

Use the following endpoints to acquire the endpoints for the Timestream for InfluxDB API.

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | timestream-influxdb.us-east-2.amazonaws.com<br>timestream-influxdb-fips.us-east-2.api.aws<br>timestream-influxdb-fips.us-east-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| US East (N. Virginia) | us-east-1 | timestream-influxdb.us-east-1.amazonaws.com<br>timestream-influxdb-fips.us-east-1.api.aws<br>timestream-influxdb-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | timestream-influxdb.us-west-2.amazonaws.com<br>timestream-influxdb-fips.us-west-2.api.aws<br>timestream-influxdb-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS | |
| Africa (Cape Town) | af-south-1 | timestream-influxdb.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | timestream-influxdb.ap-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | timestream-influxdb.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | timestream-influxdb.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | timestream-influxdb.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | timestream-influxdb.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | timestream-influxdb.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | timestream-influxdb.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | timestream-influxdb.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | timestream-influxdb.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Thailand) | ap-southeast-7 | timestream-influxdb.ap-southeast-7.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | timestream-influxdb.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | timestream-influxdb.ca-central-1.amazonaws.com<br>timestream-influxdb-fips.ca-central-1.api.aws<br>timestream-influxdb-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS<br>HTTPS |
| Europe (Frankfurt) | eu-central-1 | timestream-influxdb.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | timestream-influxdb.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | timestream-influxdb.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | timestream-influxdb.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | timestream-influxdb.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | timestream-influxdb.eu-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | timestream-influxdb.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | timestream-influxdb.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | timestream-influxdb.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | timestream-influxdb.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | timestream-influxdb.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | timestream-influxdb.sa-east-1.amazonaws.com | HTTPS | |

# Service quotas

## Timestream for Live Analytics quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CreateBatchLoadTask request rate | Each supported Region: 1 per second | No | Maximum number of CreateBatchLoadTask requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| CreateDatabase request rate | Each supported Region: 1 per second | No | Maximum number of CreateDatabase requests per second. |
| CreateScheduledQuery request rate | Each supported Region: 1 per second | No | Maximum number of CreateScheduledQuery requests per second. |
| CreateTable request rate | Each supported Region: 1 per second | No | Maximum number of CreateTable requests per second. |
| Databases per account | Each supported Region: 500 | No | The maximum number of databases you can create per AWS account. |
| DeleteDatabase request rate | Each supported Region: 1 per second | No | Maximum number of DeleteDatabase requests per second. |
| DeleteScheduledQuery request rate | Each supported Region: 1 per second | No | Maximum number of DeleteScheduledQuery requests per second. |
| DeleteTable request rate | Each supported Region: 1 per second | No | Maximum number of DeleteTable requests per second. |
| DescribeAccountSettings request rate | Each supported Region: 1 per second | No | Maximum number of DescribeAccountSettings requests per second. |
| DescribeBatchLoadTask request rate | Each supported Region: 5 per second | No | Maximum number of DescribeBatchLoadTask requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DescribeDatabase request rate | Each supported Region: 5 per second | No | Maximum number of DescribeDatabase requests per second. |
| DescribeScheduledQuery request rate | Each supported Region: 5 per second | No | Maximum number of DescribeScheduledQuery requests per second. |
| DescribeTable request rate | Each supported Region: 5 per second | No | Maximum number of DescribeTable requests per second. |
| Dimensions per table | Each supported Region: 128 | No | The maximum number of dimensions per table. |
| ExecuteScheduledQuery request rate | Each supported Region: 5 per second | No | Maximum number of ExecuteScheduledQuery requests per second. |
| ListBatchLoadTasks request rate | Each supported Region: 5 per second | No | Maximum number of ListBatchLoadTasks requests per second. |
| ListDatabases request rate | Each supported Region: 5 per second | No | Maximum number of ListDatabases requests per second. |
| ListScheduledQueries request rate | Each supported Region: 5 per second | No | Maximum number of ListScheduledQueries requests per second. |
| ListTables request rate | Each supported Region: 5 per second | No | Maximum number of ListTables requests per second. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Max allowed MaxQueryTCU | Each supported Region: 1,000 | No | The maximum allowed MaxQueryTCU in your account |
| Max allowed ProvisionedQueryTCU | Each supported Region: 1,000 | No | The maximum allowed ProvisionedQueryTCU in your account |
| Maximum count of active magnetic store partitions per database | Each supported Region: 250 | No | Maximum number of active magnetic store partitions per database. A partition may remain active for up to 6 hours after receiving ingestion. |
| Measures per table | Each supported Region: 8,192 | No | The maximum number of measures per table. |
| PrepareQuery request rate | Each supported Region: 5 per second | No | Maximum number of PrepareQuery requests per second. |
| ResumeBatchLoadTask request rate | Each supported Region: 1 per second | No | Maximum number of ResumeBatchLoadTask requests per second. |
| Scheduled queries per account | Each supported Region: 10,000 | No | The maximum number of scheduled queries you can create per AWS account. |
| Tables per account | Each supported Region: 50,000 | No | The maximum number of tables you can create per AWS account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Unique measures across multi-measure records per table | Each supported Region: 1,024 | No | The unique measures in all the multi-measure records defined in a single table. A single multi-measure record can have MaximumMeasureAttributesPerRecord unique measures, and a table can have MaximumMeasureAttributesPerTable unique measures defined across multi-measure records. |
| UpdateAccountSettings request rate | Each supported Region: 1 per second | No | Maximum number of UpdateAccountSettings requests per second. |
| UpdateDatabase request rate | Each supported Region: 1 per second | No | Maximum number of UpdateDatabase requests per second. |
| UpdateScheduledQuery request rate | Each supported Region: 1 per second | No | Maximum number of UpdateScheduledQuery requests per second. |
| UpdateTable request rate | Each supported Region: 1 per second | No | Maximum number of UpdateTable requests per second. |

For more information, see Quotas in the *Amazon Timestream Developer Guide*.

### Timestream for InfluxDB quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Database instances per account | Each supported Region: 40 | No | The maximum number of database instances you can create per AWS account. |

# AWS Telco Network Builder endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | tnb.us-east-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | tnb.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | tnb.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | tnb.ap-southeast-2.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | tnb.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | tnb.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | tnb.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | tnb.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | tnb.eu-north-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | tnb.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Concurrent ongoing network service operations | Each supported Region: 40 | Yes | The maximum number of concurrent ongoing |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | network service operations in one Region. |
| Function packages | Each supported Region: 200 | Yes | The maximum number of function packages in one Region. |
| Network packages | Each supported Region: 40 | Yes | The maximum number of network packages in one Region. |
| Network service instances | Each supported Region: 800 | Yes | The maximum number of network service instances in one Region. |

# AWS Transform endpoints and quotas

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. The following are the service quotas for this service.

## AWS Transform service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| .NET monthly lines of code (web & IDE) | Each supported Region: 2,000,000 | Yes | The maximum number of lines of .NET application code that can be transformed in a calendar month in each supported AWS Region of an AWS IAM Identity Center account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Concurrent .NET jobs with the Microsoft Visual Studio Extension per IAM Identity Center account | Each supported Region: 10 | Yes | The maximum number of concurrent .NET transformation jobs in each supported AWS Region of an IAM Identity Center account when using the Microsoft Visual Studio Extension |
| Concurrent .NET web jobs per IAM Identity Center account | Each supported Region: 5 | Yes | The maximum number of concurrent .NET transformation jobs in each supported AWS Region of an AWS IAM Identity Center account when using the AWS Transform web app. |
| Concurrent mainframe jobs per IAM Identity Center account | Each supported Region: 10 | Yes | The maximum number of concurrent Mainframe jobs in an IAM Identity Center account |
| Custom transformation definitions | Each supported Region: 1,000 | Yes | The maximum number of custom transformation definitions |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom transformation monthly AWS/dotnet-modernization agent minutes | Each supported Region: 50,000 | Yes | The maximum number of AWS/dotnet-modernization agent minutes that can be consumed in a calendar month in each supported AWS Region of an AWS IAM Identity Center account |
| Custom transformation monthly CLI conversation requests | Each supported Region: 1,000,000 | Yes | The maximum number of custom transformation conversation requests made in the CLI allowed in a calendar month |
| Workspaces per AWS IAM Identity Center | Each supported Region: 100 | Yes | The maximum number of workspaces that can be associated with an AWS IAM Identity Center |
| mainframe Analyze Lines of Code | Each supported Region: 500,000,000 | Yes | The maximum number of lines of code for Analyze Code job objective in a calendar month |
| mainframe Business Document Generation Monthly Lines of Code | Each supported Region: 250,000,000 | Yes | The maximum number of lines of code for Business Documentation Generation job objective in a calendar month |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| mainframe Decomposition Lines of Code | Each supported Region: 250,000,000 | Yes | The maximum number of lines of code for Decomposition job objective in a calendar month |
| mainframe Lines of Code per Job | ap-northeast-1: 3,000,000<br>ap-southeast-2: 3,000,000<br>eu-west-2: 3,000,000<br>Each of the other supported Regions: 20,000,000 | No | The maximum number of lines of code for each job |
| mainframe Reforge Lines of Code | Each supported Region: 50,000,000 | Yes | The maximum number of lines of code for Reforge job objective in a calendar month |
| mainframe Technical Document Generation Monthly Lines of Code | Each supported Region: 250,000,000 | Yes | The maximum number of lines of code for Technical Documentation Generation job objective in a calendar month |
| mainframe Transform of Lines of Code | Each supported Region: 250,000,000 | Yes | The maximum number of lines of code for Refactor Code job objective in a calendar month |

# Amazon Transcribe endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

Supported service endpoints differ for batch and streaming transcriptions. In addition, endpoints vary between Amazon Transcribe, Amazon Transcribe Medical, and Amazon Transcribe Call Analytics. Refer to the following tables for details.

The following table shows the supported **batch** endpoints for **Amazon Transcribe**.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | transcribe.us-east-2.amazonaws.com | HTTPS |
| | | transcribe.us-east-2.api.aws | HTTPS |
| | | fips.transcribe.us-east-2.amazonaws.com | HTTPS |
| | | transcribe-fips.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | transcribe.us-east-1.amazonaws.com | HTTPS |
| | | fips.transcribe.us-east-1.amazonaws.com | HTTPS |
| | | transcribe-fips.us-east-1.api.aws | HTTPS |
| | | transcribe.us-east-1.api.aws | HTTPS |
| US West (N. | us-west-1 | transcribe.us-west-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| California) | | transcribe-fips.us-west-1.api.aws | HTTPS | |
| | | transcribe.us-west-1.api.aws | HTTPS | |
| | | fips.transcribe.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | transcribe.us-west-2.amazonaws.com | HTTPS | |
| | | fips.transcribe.us-west-2.amazonaws.com | HTTPS | |
| | | transcribe.us-west-2.api.aws | HTTPS | |
| | | transcribe-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | transcribe.af-south-1.amazonaws.com | HTTPS | |
| | | transcribe.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | transcribe.ap-east-1.amazonaws.com | HTTPS | |
| | | transcribe.ap-east-1.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | transcribe.ap-south-1.amazonaws.com | HTTPS | |
| | | transcribe.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | transcribe.ap-northeast-2.amazonaws.com | HTTPS | |
| | | transcribe.ap-northeast-2.api.aws | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | transcribe.ap-southeast-1.amazonaws.com | HTTPS | |
| | | transcribe.ap-southeast-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | transcribe.ap-southeast-2.amazonaws.com | HTTPS |
| | | transcribe.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | transcribe.ap-northeast-1.amazonaws.com | HTTPS |
| | | transcribe.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | transcribe.ca-central-1.amazonaws.com | HTTPS |
| | | transcribe.ca-central-1.api.aws | HTTPS |
| | | fips.transcribe.ca-central-1.amazonaws.com | HTTPS |
| | | transcribe-fips.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | transcribe.eu-central-1.amazonaws.com | HTTPS |
| | | transcribe.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | transcribe.eu-west-1.amazonaws.com | HTTPS |
| | | transcribe.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | transcribe.eu-west-2.amazonaws.com | HTTPS |
| | | transcribe.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | transcribe.eu-west-3.amazonaws.com | HTTPS |
| | | transcribe.eu-west-3.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | transcribe.eu-north-1.amazonaws.com | HTTPS |
| | | transcribe.eu-north-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Zurich) | eu-central-2 | transcribe.eu-central-2.amazonaws.com | HTTPS |
| | | transcribe.eu-central-2.api.aws | HTTPS |
| Middle East (Bahrain) | me-south-1 | transcribe.me-south-1.amazonaws.com | HTTPS |
| | | transcribe.me-south-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | transcribe.sa-east-1.amazonaws.com | HTTPS |
| | | transcribe.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | transcribe.us-gov-east-1.amazonaws.com | HTTPS |
| | | transcribe.us-gov-east-1.api.aws | HTTPS |
| | | fips.transcribe.us-gov-east-1.amazonaws.com | HTTPS |
| | | transcribe-fips.us-gov-east-1.api.aws | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | transcribe.us-gov-west-1.amazonaws.com | HTTPS |
| | | transcribe-fips.us-gov-west-1.api.aws | HTTPS |
| | | fips.transcribe.us-gov-west-1.amazonaws.com | HTTPS |
| | | transcribe.us-gov-west-1.api.aws | HTTPS |

The following table shows the supported **streaming** endpoints for **Amazon Transcribe**.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | transcribestreaming.us-east-2.amazonaws.com | HTTPS |
| | | | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| | | transcribestreaming.us-east-2.api.aws | HTTPS | |
| | | transcribestreaming-fips.us-east-2.amazonaws.com | HTTPS | |
| | | transcribestreaming-fips.us-east-2.api.aws | | |
| US East (N. Virginia) | us-east-1 | transcribestreaming.us-east-1.amazonaws.com | HTTPS | |
| | | transcribestreaming-fips.us-east-1.amazonaws.com | HTTPS | |
| | | transcribestreaming-fips.us-east-1.api.aws | HTTPS | |
| | | transcribestreaming.us-east-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | transcribestreaming.us-west-2.amazonaws.com | HTTPS | |
| | | transcribestreaming-fips.us-west-2.amazonaws.com | HTTPS | |
| | | transcribestreaming.us-west-2.api.aws | HTTPS | |
| | | transcribestreaming-fips.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | transcribestreaming.af-south-1.amazonaws.com | HTTPS | |
| | | transcribestreaming.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | transcribestreaming.ap-southeast-5.amazonaws.com | HTTPS | |
| | | transcribestreaming.ap-southeast-5.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | transcribestreaming.ap-south-1.amazonaws.com<br><br>transcribestreaming.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | transcribestreaming.ap-northeast-2.amazonaws.com<br><br>transcribestreaming.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | transcribestreaming.ap-southeast-1.amazonaws.com<br><br>transcribestreaming.ap-southeast-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | transcribestreaming.ap-southeast-2.amazonaws.com<br><br>transcribestreaming.ap-southeast-2.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | transcribestreaming.ap-southeast-7.amazonaws.com<br><br>transcribestreaming.ap-southeast-7.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | transcribestreaming.ap-northeast-1.amazonaws.com<br><br>transcribestreaming.ap-northeast-1.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Canada (Central) | ca-central-1 | transcribestreaming.ca-central-1.amazonaws.com<br>transcribestreaming.ca-central-1.api.aws<br>transcribestreaming-fips.ca-central-1.amazonaws.com<br>transcribestreaming-fips.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| Europe (Frankfurt) | eu-central-1 | transcribestreaming.eu-central-1.amazonaws.com<br>transcribestreaming.eu-central-1.api.aws | HTTPS<br>HTTPS | |
| Europe (Ireland) | eu-west-1 | transcribestreaming.eu-west-1.amazonaws.com<br>transcribestreaming.eu-west-1.api.aws | HTTPS<br>HTTPS | |
| Europe (London) | eu-west-2 | transcribestreaming.eu-west-2.amazonaws.com<br>transcribestreaming.eu-west-2.api.aws | HTTPS<br>HTTPS | |
| Europe (Zurich) | eu-central-2 | transcribestreaming.eu-central-2.amazonaws.com<br>transcribestreaming.eu-central-2.api.aws | HTTPS<br>HTTPS | |
| Mexico (Central) | mx-central-1 | transcribestreaming.mx-central-1.amazonaws.com<br>transcribestreaming.mx-central-1.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| South America (São Paulo) | sa-east-1 | transcribestreaming.sa-east-1.amazonaws.com<br>transcribestreaming.sa-east-1.api.aws | HTTPS<br>HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | transcribestreaming.us-gov-east-1.amazonaws.com<br>transcribestreaming.us-gov-east-1.api.aws<br>transcribestreaming-fips.us-gov-east-1.amazonaws.com<br>transcribestreaming-fips.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |
| AWS GovCloud (US-West) | us-gov-west-1 | transcribestreaming.us-gov-west-1.amazonaws.com<br>transcribestreaming-fips.us-gov-west-1.api.aws<br>transcribestreaming-fips.us-gov-west-1.amazonaws.com<br>transcribestreaming.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

The following tables show the supported AWS Regions for **Amazon Transcribe Medical** and **Amazon Transcribe Call Analytics**. Note that 'post-call' refers to a batch transcription and 'real-time' refers to a streaming transcription.

To get the endpoints for these Regions, refer to the preceding Amazon Transcribe endpoints tables.

### Amazon Transcribe Medical

| Region | Transcription type |
| --- | --- |
| ap-east-1 (Hong Kong) | batch |
| ap-northeast-1 (Tokyo) | batch, streaming |
| ap-northeast-2 (Seoul) | batch, streaming |
| ap-south-1 (Mumbai) | batch |
| ap-southeast-1 (Singapore) | batch |
| ap-southeast-2 (Sydney) | batch, streaming |
| ca-central-1 (Canada, Central) | batch, streaming |
| eu-central-1 (Frankfurt) | batch, streaming |
| eu-west-1 (Ireland) | batch, streaming |
| eu-west-2 (London) | batch, streaming |
| eu-west-3 (Paris) | batch |
| me-south-1 (Bahrain) | batch |
| sa-east-1 (São Paulo) | batch, streaming |
| us-east-1 (N. Virginia) | batch, streaming |
| us-east-2 (Ohio) | batch, streaming |
| us-gov-east-1 (GovCloud, US-East) | batch, streaming |
| us-gov-west-1 (GovCloud, US-West) | batch, streaming |
| us-west-1 (San Francisco) | batch |
| us-west-2 (Oregon) | batch, streaming |
| eu-north-1 (Stockholm) | batch |

| Region | Transcription type |
| --- | --- |
| af-south-1 (Cape Town) | batch |

### Amazon Transcribe Call Analytics

| Region | Transcription type |
| --- | --- |
| ap-northeast-1 (Tokyo) | post-call, real-time |
| ap-northeast-2 (Seoul) | post-call, real-time |
| ap-south-1 (Mumbai) | post-call |
| ap-southeast-1 (Singapore) | post-call |
| ap-southeast-2 (Sydney) | post-call, real-time |
| ca-central-1 (Canada, Central) | post-call, real-time |
| eu-central-1 (Frankfurt) | post-call, real-time |
| eu-west-2 (London) | post-call, real-time |
| us-east-1 (N. Virginia) | post-call, real-time |
| us-west-2 (Oregon) | post-call, real-time |

## Service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Channel identification limit | Each supported Region: 2 | Yes | The maximum number of audio channels that can be identified and extracted from a media file. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom medical vocabulary size limit | Each supported Region: 51,200 | Yes | The maximum size of custom medical vocabulary in this account in the current region in bytes. |
| Custom vocabulary size limit | Each supported Region: 51,200 | Yes | The maximum size of custom vocabulary in this account in the current region in bytes. |
| Job queue bandwidth ratio | Each supported Region: 90 | Yes | The maximum job queue bandwidth ratio in this account in the current region. |
| Maximum audio file length | Each supported Region: 28,800 Seconds | No | The maximum length of an audio file in seconds. |
| Maximum audio file length (Medical) | Each supported Region: 14,400 Seconds | No | The maximum audio file length in seconds (Medical). |
| Maximum audio file length for Call Analytics batch jobs | Each supported Region: 14,400 Seconds | No | The maximum audio file length (in seconds) for Call Analytics batch jobs. |
| Maximum audio file length for MedicalScribe jobs | Each supported Region: 5,400 Seconds | No | The maximum audio file length (in seconds) for MedicalScribe jobs. |
| Maximum audio file size | Each supported Region: 2 Gigabytes | No | The maximum size of an audio file in Gigabytes. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum audio file size (Medical) | Each supported Region: 2 Gigabytes | No | The maximum audio file size (in GB) for medical transcription. |
| Maximum audio file size for Call Analytics batch jobs | Each supported Region: 500 Megabytes | No | The maximum audio file size (in MB) for Call Analytics batch jobs. |
| Maximum audio file size for MedicalScribe jobs | Each supported Region: 2 Gigabytes | No | The maximum audio file size (in GB) for MedicalScribe jobs. |
| Maximum length of a custom vocabulary phrase | Each supported Region: 256 | No | The maximum number of characters in a custom vocabulary phrase. |
| Maximum number of categories for Call Analytics batch jobs | Each supported Region: 200 | Yes | The maximum number of categories per account for Call Analytics batch jobs. |
| Maximum number of rules per category for Call Analytics batch jobs | Each supported Region: 20 | Yes | The maximum number of rules per category for Call Analytics batch jobs. |
| Maximum number of targets allowed per category for Call Analytics batch jobs | Each supported Region: 100 | Yes | The maximum number of targets allowed per category for Call Analytics batch jobs |
| Maximum number of vocabulary filters | Each supported Region: 100 | No | The maximum number of vocabulary filters from this account in the current Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum size of a vocabulary filter | Each supported Region: 50 Kilobytes | No | The maximum size (in KB) of a vocabulary filter. |
| Minimum audio file duration | Each supported Region: 500 Milliseconds | No | The minimum audio file duration (in ms). |
| Minimum audio file duration (Medical) | Each supported Region: 500 Milliseconds | No | The minimum audio file duration, in milliseconds (ms) for medical transcription. |
| Minimum audio file duration for Call Analytics batch jobs | Each supported Region: 500 Milliseconds | No | The minimum audio file duration, in milliseconds (ms), for Call Analytics batch jobs. |
| Number of channels for channel identification (Medical) | Each supported Region: 2 | No | The maximum number of channels for channel identification (Medical). |
| Number of channels for channel identification for Call Analytics batch jobs | Each supported Region: 2 | No | The maximum number of channels for channel identification for Call Analytics batch jobs. |
| Number of concurrent Call Analytics streams (HTTP/2 + Websocket) | Each supported Region: 25 | Yes | The maximum number of concurrent Call Analytics stream transcription sessions in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of concurrent Medical streams (HTTP/2 + Websocket) | Each supported Region: 25 | Yes | The maximum number of concurrent Medical stream transcription sessions in this account in the current region. |
| Number of concurrent MedicalScribe post-stream analytics jobs | Each supported Region: 1 | Yes | The maximum number of concurrent MedicalScribe post-stream analytics jobs in this account in the current region. |
| Number of concurrent MedicalScribe streams (HTTP/2) | Each supported Region: 10 | Yes | The maximum number of concurrent MedicalScribe stream sessions in this account in the current region. |
| Number of concurrent analytics jobs | Each supported Region: 100 | Yes | The maximum number of concurrent analytics jobs in this account in the current region. |
| Number of concurrent language models | Each supported Region: 3 | Yes | The maximum number of concurrent language models in this account in the current region. |
| Number of concurrent medical notes jobs | Each supported Region: 5 | Yes | The maximum number of concurrent medical notes jobs in this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Number of concurrent medical transcription jobs | Each supported Region: 250 | Yes | The maximum number of concurrent medical transcription jobs in this account in the current region. |
| Number of concurrent medical vocabulary jobs | Each supported Region: 10 | Yes | The maximum number of concurrent medical vocabulary jobs in this account in the current region. |
| Number of concurrent streams (HTTP/2 + Websocket) | Each supported Region: 25 | Yes | The maximum number of concurrent stream transcription sessions in this account in the current region. |
| Number of concurrent transcription jobs | Each supported Region: 250 | Yes | The maximum number of concurrent transcription jobs in this account in the current region. |
| Number of concurrent vocabulary jobs | Each supported Region: 10 | Yes | The maximum number of concurrent vocabulary jobs in this account in the current region. |
| Number of days that job records are retained | Each supported Region: 90 | No | The number of days that job records are retained. |
| Number of days that job records are retained (Medical) | Each supported Region: 90 | No | The maximum number of days that job records are retained (Medical). |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| Number of days that job records are retained for Call Analytics batch jobs | Each supported Region: 90 | No | The number of days job records are retained for Call Analytics batch jobs. |
| Total language model limit | Each supported Region: 10 | Yes | The maximum total number of language models in this account in the current region. |
| Total medical vocabulary limit | Each supported Region: 100 | Yes | The maximum total number of medical vocabularies in this account in the current region. |
| Total vocabulary limit | Each supported Region: 100 | Yes | The maximum total number of vocabularies in this account in the current region. |
| Transactions per second, CreateCallAnalyticsCategory | Each supported Region: 10 per second | Yes | The maximum number of CreateCallAnalyticsCategory requests that you can make per second from this account in the current region. |
| Transactions per second, CreateLanguageModel | Each supported Region: 5 per second | Yes | The maximum number of CreateLanguageModel requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, CreateMedicalVocabulary | Each supported Region: 10 per second | Yes | The maximum number of CreateMedicalVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, CreateVocabulary | Each supported Region: 10 per second | Yes | The maximum number of CreateVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, CreateVocabularyFilter | Each supported Region: 10 per second | Yes | The maximum number of CreateVocabularyFilter requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteCallAnalyticsCategory | Each supported Region: 5 per second | Yes | The maximum number of DeleteCallAnalyticsCategory requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteCallAnalyticsJob | Each supported Region: 5 per second | Yes | The maximum number of DeleteCallAnalyticsJob requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second, DeleteLanguageModel | Each supported Region: 5 per second | Yes | The maximum number of DeleteLanguageModel requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteMedicalScribeJob | Each supported Region: 5 per second | Yes | The maximum number of DeleteMedicalScribeJob requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteMedicalTranscriptionJob | Each supported Region: 5 per second | Yes | The maximum number of DeleteMedicalTranscriptionJob requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteMedicalVocabulary | Each supported Region: 5 per second | Yes | The maximum number of DeleteMedicalVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteTranscriptionJob | Each supported Region: 5 per second | Yes | The maximum number of DeleteTranscriptionJob requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second, DeleteVocabulary | Each supported Region: 5 per second | Yes | The maximum number of DeleteVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, DeleteVocabularyFilter | Each supported Region: 5 per second | Yes | The maximum number of DeleteVocabularyFilter requests that you can make per second from this account in the current region. |
| Transactions per second, DescribeLanguageModel | Each supported Region: 10 per second | Yes | The maximum number of DescribeLanguageModel requests that you can make per second from this account in the current region. |
| Transactions per second, GetCallAnalyticsCategory | Each supported Region: 20 per second | Yes | The maximum number of GetCallAnalyticsCategory requests that you can make per second from this account in the current region. |
| Transactions per second, GetCallAnalyticsJob | Each supported Region: 20 per second | Yes | The maximum number of GetCallAnalyticsJob requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, GetMedicalScribeJob | Each supported Region: 5 per second | Yes | The maximum number of GetMedicalScribeJob requests that you can make per second from this account in the current region. |
| Transactions per second, GetMedicalScribeStream | Each supported Region: 10 per second | Yes | The maximum number of GetMedicalScribeStream requests that you can make per second from this account in the current region. |
| Transactions per second, GetMedicalTranscriptionJob | Each supported Region: 30 per second | Yes | The maximum number of GetMedicalTranscriptionJob requests that you can make per second from this account in the current region. |
| Transactions per second, GetMedicalVocabulary | Each supported Region: 20 per second | Yes | The maximum number of GetMedicalVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, GetTranscriptionJob | Each supported Region: 30 per second | Yes | The maximum number of GetTranscriptionJob requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Transactions per second, GetVocabulary | Each supported Region: 20 per second | Yes | The maximum number of GetVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, GetVocabularyFilter | Each supported Region: 20 per second | Yes | The maximum number of GetVocabularyFilter requests that you can make per second from this account in the current region. |
| Transactions per second, ListCallAnalyticsCategories | Each supported Region: 5 per second | Yes | The maximum number of ListCallAnalyticsCategories requests that you can make per second from this account in the current region. |
| Transactions per second, ListCallAnalyticsJobs | Each supported Region: 5 per second | Yes | The maximum number of ListCallAnalyticsJobs requests that you can make per second from this account in the current region. |
| Transactions per second, ListLanguageModels | Each supported Region: 10 per second | Yes | The maximum number of ListLanguageModels requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, ListMedicalScribeJobs | Each supported Region: 5 per second | Yes | The maximum number of ListMedicalScribeJobs requests that you can make per second from this account in the current region. |
| Transactions per second, ListMedicalTranscriptionJobs | Each supported Region: 5 per second | Yes | The maximum number of ListMedicalTranscriptionJobs requests that you can make per second from this account in the current region. |
| Transactions per second, ListMedicalVocabularies | Each supported Region: 5 per second | Yes | The maximum number of ListMedicalVocabularies requests that you can make per second from this account in the current region. |
| Transactions per second, ListTagsForResource | Each supported Region: 50 per second | Yes | The maximum number of ListTagsForResource requests that you can make per second from this account in the current region. |
| Transactions per second, ListTranscriptionJobs | Each supported Region: 5 per second | Yes | The maximum number of ListTranscriptionJobs requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, ListVocabularies | Each supported Region: 5 per second | Yes | The maximum number of ListVocabularies requests that you can make per second from this account in the current region. |
| Transactions per second, ListVocabularyFilters | Each supported Region: 5 per second | Yes | The maximum number of ListVocabularyFilters requests that you can make per second from this account in the current region. |
| Transactions per second, StartCallAnalyticsJob | Each supported Region: 10 per second | Yes | The maximum number of StartCallAnalyticsJob requests that you can make per second from this account in the current region. |
| Transactions per second, StartCallAnalyticsStreamTranscription | Each supported Region: 25 per second | Yes | The maximum number of StartCallAnalyticsStreamTranscription requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, StartCallAnalyticsStreamTranscriptionWebsocket | Each supported Region: 25 per second | Yes | The maximum number of StartCallAnalyticsStreamTranscriptionWebsocket requests that you can make per second from this account in the current region. |
| Transactions per second, StartMedicalScribeJob | Each supported Region: 5 per second | Yes | The maximum number of StartMedicalScribeJob requests that you can make per second from this account in the current region. |
| Transactions per second, StartMedicalScribeStream | Each supported Region: 2 per second | Yes | The maximum number of StartMedicalScribeStream requests that you can make per second from this account in the current region. |
| Transactions per second, StartMedicalStreamTranscription | Each supported Region: 25 per second | Yes | The maximum number of StartMedicalStreamTranscription requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, StartMedicalStreamTranscriptionWebsocket | Each supported Region: 25 per second | Yes | The maximum number of StartMedicalStreamTranscriptionWebsocket requests that you can make per second from this account in the current region. |
| Transactions per second, StartMedicalTranscriptionJob | Each supported Region: 25 per second | Yes | The maximum number of StartMedicalTranscriptionJob requests that you can make per second from this account in the current region. |
| Transactions per second, StartStreamTranscription | Each supported Region: 25 per second | Yes | The maximum number of StartStreamTranscription requests that you can make per second from this account in the current region. |
| Transactions per second, StartStreamTranscriptionWebSocket | Each supported Region: 25 per second | Yes | The maximum number of StartStreamTranscriptionWebsocket requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, StartTranscriptionJob | Each supported Region: 25 per second | Yes | The maximum number of StartTranscriptionJob requests that you can make per second from this account in the current region. |
| Transactions per second, TagResource | Each supported Region: 10 per second | Yes | The maximum number of TagResource requests that you can make per second from this account in the current region. |
| Transactions per second, UntagResource | Each supported Region: 10 per second | Yes | The maximum number of UntagResource requests that you can make per second from this account in the current region. |
| Transactions per second, UpdateCallAnalyticsCategory | Each supported Region: 10 per second | Yes | The maximum number of UpdateCallAnalyticsCategory requests that you can make per second from this account in the current region. |
| Transactions per second, UpdateLanguageModel | Each supported Region: 5 per second | Yes | The maximum number of UpdateLanguageModel requests that you can make per second from this account in the current region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Transactions per second, UpdateMedicalVocabulary | Each supported Region: 10 per second | Yes | The maximum number of UpdateMedicalVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, UpdateVocabulary | Each supported Region: 10 per second | Yes | The maximum number of UpdateVocabulary requests that you can make per second from this account in the current region. |
| Transactions per second, UpdateVocabularyFilter | Each supported Region: 10 per second | Yes | The maximum number of UpdateVocabularyFilter requests that you can make per second from this account in the current region. |

You can view your current quotas in the AWS Management Console.

## AWS Transfer Family endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | transfer.us-east-2.amazonaws.com | HTTPS |
| | | transfer-fips.us-east-2.amazonaws.com | HTTPS |
| | | transfer.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | transfer.us-east-1.amazonaws.com | HTTPS |
| | | transfer-fips.us-east-1.amazonaws.com | HTTPS |
| | | transfer.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | transfer.us-west-1.amazonaws.com | HTTPS |
| | | transfer-fips.us-west-1.amazonaws.com | HTTPS |
| | | transfer.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | transfer.us-west-2.amazonaws.com | HTTPS |
| | | transfer-fips.us-west-2.amazonaws.com | HTTPS |
| | | transfer.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | transfer.af-south-1.amazonaws.com | HTTPS |
| | | transfer.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | transfer.ap-east-1.amazonaws.com | HTTPS |
| | | transfer.ap-east-1.api.aws | HTTPS |
| Asia Pacific | ap-south-2 | transfer.ap-south-2.amazonaws.com | HTTPS |
| | | transfer.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | transfer.ap-southeast-3.amazonaws.com<br>transfer.ap-southeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | transfer.ap-southeast-5.amazonaws.com<br>transfer.ap-southeast-5.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | transfer.ap-southeast-4.amazonaws.com<br>transfer.ap-southeast-4.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | transfer.ap-south-1.amazonaws.com<br>transfer.ap-south-1.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | transfer.ap-southeast-6.amazonaws.com<br>transfer.ap-southeast-6.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | transfer.ap-northeast-3.amazonaws.com<br>transfer.ap-northeast-3.api.aws | HTTPS<br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | transfer.ap-northeast-2.amazonaws.com<br>transfer.ap-northeast-2.api.aws | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | transfer.ap-southeast-1.amazonaws.com | HTTPS |
| | | transfer.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | transfer.ap-southeast-2.amazonaws.com | HTTPS |
| | | transfer.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | transfer.ap-east-2.amazonaws.com | HTTPS |
| | | transfer.ap-east-2.api.aws | HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | transfer.ap-southeast-7.amazonaws.com | HTTPS |
| | | transfer.ap-southeast-7.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | transfer.ap-northeast-1.amazonaws.com | HTTPS |
| | | transfer.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | transfer.ca-central-1.amazonaws.com | HTTPS |
| | | transfer-fips.ca-central-1.amazonaws.com | HTTPS |
| | | transfer.ca-central-1.api.aws | HTTPS |
| Canada West (Calgary) | ca-west-1 | transfer.ca-west-1.amazonaws.com | HTTPS |
| | | transfer-fips.ca-west-1.amazonaws.com | HTTPS |
| | | transfer.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | transfer.eu-central-1.amazonaws.com | HTTPS |
| | | transfer.eu-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Ireland) | eu-west-1 | transfer.eu-west-1.amazonaws.com | HTTPS | |
| | | transfer.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | transfer.eu-west-2.amazonaws.com | HTTPS | |
| | | transfer.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | transfer.eu-south-1.amazonaws.com | HTTPS | |
| | | transfer.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | transfer.eu-west-3.amazonaws.com | HTTPS | |
| | | transfer.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | transfer.eu-south-2.amazonaws.com | HTTPS | |
| | | transfer.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | transfer.eu-north-1.amazonaws.com | HTTPS | |
| | | transfer.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | transfer.eu-central-2.amazonaws.com | HTTPS | |
| | | transfer.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | transfer.il-central-1.amazonaws.com | HTTPS | |
| | | transfer.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | transfer.mx-central-1.amazonaws.com | HTTPS | |
| | | transfer.mx-central-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Middle East (Bahrain) | me-south-1 | transfer.me-south-1.amazonaws.com<br>transfer.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | transfer.me-central-1.amazonaws.com<br>transfer.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | transfer.sa-east-1.amazonaws.com<br>transfer.sa-east-1.api.aws | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | transfer.us-gov-east-1.amazonaws.com<br>transfer-fips.us-gov-east-1.amazonaws.com<br>transfer.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | transfer.us-gov-west-1.amazonaws.com<br>transfer-fips.us-gov-west-1.amazonaws.com<br>transfer.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| AS2 connectors per account | Each supported Region: 100 | Yes | Maximum number of AS2 connectors per account |
| Agreements per account | Each supported Region: 100 | Yes | Maximum number of agreements per account |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Certificates per account | Each supported Region: 1,000 | Yes | Maximum number of certificates per account |
| Certificates per profile | Each supported Region: 10 | No | Maximum number of certificates per profile |
| Concurrent AS2 messages per connector | Each supported Region: 400 | No | The maximum number of concurrent AS2 messages that can be handled by a connector |
| Concurrent AS2 messages per server | Each supported Region: 400 | No | The maximum number of concurrent AS2 messages that can be handled by a server |
| Concurrent multiplexed SFTP sessions per connection | Each supported Region: 10 | No | The maximum number of concurrent multiplexed SFTP sessions per connection |
| Concurrent sessions per server | Each supported Region: 10,000 | No | Maximum number of concurrent sessions per server |
| File size | Each supported Region: 5 Terabytes | No | Maximum size of an individual file, which is the maximum size of an individual object in Amazon S3 |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Idle connection timeout | Each supported Region: 1,800 Seconds | No | Inactivity timeout for SFTP/FTP(S) connections. If there is no activity after the period has passed the client may be disconnected |
| Maximum inbound AS2 message size | Each supported Region: 1,000 Megabytes | No | Maximum size in (MB) of an inbound AS2 message |
| Maximum logical directory mappings characters | Each supported Region: 2,100,000 Bytes | No | The maximum length in UTF-8 encoded characters of logical directory mappings JSON document per user |
| Maximum logical directory mappings entries | Each supported Region: 2,000 | Yes | The maximum entries per logical directory mapping per Service Managed user |
| Maximum number of AD Groups for access | Each supported Region: 100 | Yes | Number of Active Directory Groups allowed for mapping to access per server |
| Maximum number of new executions per workflow | Each supported Region: 100 | No | Maximum number of new executions allowed per workflow at one time. |
| Maximum outbound AS2 message size | Each supported Region: 1,000 Megabytes | No | Maximum size in (MB) of an outbound AS2 message |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum web app units per web app | Each supported Region: 4 | Yes | The maximum number of web app units per web app resource |
| Maximum web apps per account | Each supported Region: 10 | Yes | The maximum number of web apps per account |
| New executions refill rate per workflow per second | Each supported Region: 1 | No | The new executions refill rate per workflow per second |
| Number of Service Managed users per server | Each supported Region: 10,000 | Yes | Maximum number of Service Managed users per server |
| Number of authentication requests per user per second | Each supported Region: 2 | No | If your identity provider type is AWS Directory Service, there is a limit of 2 authentications per user per second per directory |
| Number of files per StartFileTransfer request | Each supported Region: 10 per second | No | Maximum number of files per StartFileTransfer request |
| Profiles per account | Each supported Region: 1,000 | Yes | Maximum number of profiles per account |
| Rate of AS2 messages per server | Each supported Region: 100 per second | No | The maximum number of AS2 messages a server can receive per second |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of StartDirectoryListing per SFTP connector | Each supported Region: 3 per second | No | Maximum number of StartDirectoryListing request sent per second per SFTP connector |
| Rate of StartFileTransfer files per AS2 connector | Each supported Region: 100 per second | Yes | Maximum number of StartFileTransfer files sent per second per AS2 connector |
| Rate of StartFileTransfer files per SFTP connector | Each supported Region: 100 per second | Yes | Maximum number of StartFileTransfer files sent per second per SFTP connector |
| Rate of StartRemoteDelete files per SFTP connector | Each supported Region: 100 per second | Yes | Maximum number of StartRemoteDelete files sent per second per SFTP connector |
| Rate of StartRemoteMove files per SFTP connector | Each supported Region: 100 per second | Yes | Maximum number of StartRemoteMove files sent per second per SFTP connector |
| SFTP connectors per account | Each supported Region: 100 | Yes | Maximum number of SFTP connectors per account |
| SSH keys per Service Managed user | Each supported Region: 50 | Yes | Maximum number of SSH keys per Service Managed user |
| Servers per account | Each supported Region: 50 | Yes | Maximum number of servers per account |

| Name | Default | Adjuse | Description |
| --- | --- | --- | --- |
| VPC_ENDPOINT servers per account | Each supported Region: 10 | No | Maximum number of VPC_ENDPOINT servers per account |
| Workflows per account | Each supported Region: 10 | Yes | Maximum number of workflows per account |

# Amazon Translate endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | translate.us-east-2.amazonaws.com | HTTPS | |
| | | translate.us-east-2.api.aws | HTTPS | |
| | | translate-fips.us-east-2.api.aws | HTTPS | |
| | | translate-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | translate.us-east-1.amazonaws.com | HTTPS | |
| | | translate-fips.us-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | translate-fips.us-east-1.api.aws | HTTPS |
| | | translate.us-east-1.api.aws | HTTPS |
| US West (N. California) | us-west-1 | translate.us-west-1.amazonaws.com | HTTPS |
| | | translate-fips.us-west-1.api.aws | HTTPS |
| | | translate.us-west-1.api.aws | HTTPS |
| | | translate-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | translate.us-west-2.amazonaws.com | HTTPS |
| | | translate-fips.us-west-2.amazonaws.com | HTTPS |
| | | translate.us-west-2.api.aws | HTTPS |
| | | translate-fips.us-west-2.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | translate.ap-east-1.amazonaws.com | HTTPS |
| | | translate.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | translate.ap-south-1.amazonaws.com | HTTPS |
| | | translate.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | translate.ap-northeast-2.amazonaws.com | HTTPS |
| | | translate.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | translate.ap-southeast-1.amazonaws.com | HTTPS |
| | | translate.ap-southeast-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | translate.ap-southeast-2.amazonaws.com | HTTPS | |
| | | translate.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | translate.ap-northeast-1.amazonaws.com | HTTPS | |
| | | translate.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | translate.ca-central-1.amazonaws.com | HTTPS | |
| | | translate.ca-central-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | translate.eu-central-1.amazonaws.com | HTTPS | |
| | | translate.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | translate.eu-west-1.amazonaws.com | HTTPS | |
| | | translate.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | translate.eu-west-2.amazonaws.com | HTTPS | |
| | | translate.eu-west-2.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | translate.eu-west-3.amazonaws.com | HTTPS | |
| | | translate.eu-west-3.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | translate.eu-north-1.amazonaws.com | HTTPS | |
| | | translate.eu-north-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | translate.us-gov-west-1.amazonaws.com | HTTPS | |
| | | translate-fips.us-gov-west-1.api.aws | HTTPS | |
| | | translate.us-gov-west-1.api.aws | HTTPS | |
| | | translate-fips.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjuse | Description |
|---|---|---|---|
| Concurrent batch translation jobs | Each supported Region: 10 | Yes | The maximum number of concurrent batch translation jobs in this account in the current Region. |
| Custom terminology files | Each supported Region: 100 | Yes | The maximum number of custom terminology files you can store in this account in the current Region. |
| Parallel data resources | Each supported Region: 1,000 | Yes | The maximum number of parallel data resources in this account in the current Region. |

For more information, see Guidelines and Quotas in the *Amazon Translate Developer Guide*.

# AWS Trusted Advisor API endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | trustedadvisor.us-east-1.api.aws | HTTPS |
| US East (Ohio) | us-east-2 | trustedadvisor.us-east-2.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | trustedadvisor.us-west-2.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | trustedadvisor.eu-west-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | trustedadvisor.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | trustedadvisor.ap-southeast-2.api.aws | HTTPS |

# AWS User Notifications endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | notifications.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | notifications.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | notifications.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | notifications.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | notifications.af-south-1.api.aws | HTTPS | |
| Asia Pacific | ap-east-1 | notifications.ap-east-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hong Kong) | | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | notifications.ap-south-2.api.aws | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | notifications.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | notifications.ap-southeast-5.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | notifications.ap-southeast-4.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | notifications.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | notifications.ap-northeast-3.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | notifications.ap-northeast-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | notifications.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | notifications.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | notifications.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | notifications.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | notifications.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | notifications.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | notifications.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | notifications.eu-west-1.api.aws | HTTPS | |
| Europe (London) | eu-west-2 | notifications.eu-west-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Milan) | eu-south-1 | notifications.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | notifications.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | notifications.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | notifications.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | notifications.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | notifications.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | notifications.mx-central-1.api.aws | HTTPS | |
| Middle East (Bahrain) | me-south-1 | notifications.me-south-1.api.aws | HTTPS | |
| Middle East (UAE) | me-central-1 | notifications.me-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | notifications.sa-east-1.api.aws | HTTPS | |

# AWS User Notifications Contacts endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | notifications-contacts.us-east-1.api.aws | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Notification configurations total for an AWS account | 50 notification configurations. | No | The maximum number of notification configurations that you can create in an AWS account. |
| Notification configurations for a single Service | 20 notification configurations for any specific service for an AWS account. | No | The maximum number of notification configurations that you can create for a given service in an AWS account. |
| Notification configurations per Service and Event type | 10 notification configurations for each service and event type for an AWS account. | No | The maximum number of notification configurations by Service and Event type you can create for a given AWS account. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Event rules for a given notification configuration | 10 event rules | No | The maximum number of event rules that you can create for each notification configuration in your AWS account. |
| Channels for a given notification configuration | 50 channels (email, mobile devices, or chat channels) for each notification configuration. | No | The maximum number of channels for each notification configuration that you can create in your AWS account. |
| Email contacts | 500 email contacts for each AWS account. | No | The maximum number of email contacts that you can add for each AWS account. |
| Notification hubs | 3 hubs for each AWS account. | No | The maximum number of notification hubs you can add to each AWS account. |
| Rate of source events for a given AWS account | 1 per second. | No | The maximum number of source events per second you can receive in each AWS account. |

# AWS User Experience Customization endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | uxc.us-east-2.api.aws | HTTPS | |
| US East (N. Virginia) | us-east-1 | uxc.us-east-1.api.aws | HTTPS | |
| US West (N. California) | us-west-1 | uxc.us-west-1.api.aws | HTTPS | |
| US West (Oregon) | us-west-2 | uxc.us-west-2.api.aws | HTTPS | |
| Africa (Cape Town) | af-south-1 | uxc.af-south-1.api.aws | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | uxc.ap-east-1.api.aws | HTTPS | |
| Asia Pacific | ap-south-2 | uxc.ap-south-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | uxc.ap-southeast-3.api.aws | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | uxc.ap-southeast-5.api.aws | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | uxc.ap-southeast-4.api.aws | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | uxc.ap-south-1.api.aws | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | uxc.ap-southeast-6.api.aws | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | uxc.ap-northeast-3.api.aws | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | uxc.ap-northeast-2.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | uxc.ap-southeast-1.api.aws | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | uxc.ap-southeast-2.api.aws | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | uxc.ap-east-2.api.aws | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | uxc.ap-southeast-7.api.aws | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | uxc.ap-northeast-1.api.aws | HTTPS | |
| Canada (Central) | ca-central-1 | uxc.ca-central-1.api.aws | HTTPS | |
| Canada West (Calgary) | ca-west-1 | uxc.ca-west-1.api.aws | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | uxc.eu-central-1.api.aws | HTTPS | |
| Europe (Ireland) | eu-west-1 | uxc.eu-west-1.api.aws | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | uxc.eu-west-2.api.aws | HTTPS | |
| Europe (Milan) | eu-south-1 | uxc.eu-south-1.api.aws | HTTPS | |
| Europe (Paris) | eu-west-3 | uxc.eu-west-3.api.aws | HTTPS | |
| Europe (Spain) | eu-south-2 | uxc.eu-south-2.api.aws | HTTPS | |
| Europe (Stockholm) | eu-north-1 | uxc.eu-north-1.api.aws | HTTPS | |
| Europe (Zurich) | eu-central-2 | uxc.eu-central-2.api.aws | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | uxc.il-central-1.api.aws | HTTPS | |
| Mexico (Central) | mx-central-1 | uxc.mx-central-1.api.aws | HTTPS | |
| South America (São Paulo) | sa-east-1 | uxc.sa-east-1.api.aws | HTTPS | |

## Service quotas

This service has no quotas.

# AWS Verified Access endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The API actions to manage AWS Verified Access resources (for example, Verified Access endpoints, Verified Access groups, and Verified Access trust providers) are part of the Amazon EC2 API. For more information, see AWS Verified Access actions in the *Amazon EC2 API Reference*.

For the service endpoints for Amazon EC2, see the section called “Amazon EC2”.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Verified Access Instances | 5 | Yes | The maximum number of Verified Access Instances that customers can create in the current Region. |
| Verified Access Groups | 10 | Yes | The maximum number of Verified Access Groups that customers can create in the current Region. |
| Verified Access Trust Providers | 15 | Yes | The maximum number of Verified Access Trust Providers that customers can create in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Verified Access Endpoints | 50 | Yes | The maximum number of Verified Access Endpoints that customers can create in the current Region. |

# Amazon Verified Permissions endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
| --- | --- | --- | --- |
| US East (Ohio) | us-east-2 | verifiedpermissions.us-east-2.amazonaws.com | HTTPS |
| | | verifiedpermissions-fips.us-east-2.api.aws | HTTPS |
| | | verifiedpermissions-fips.us-east-2.amazonaws.com | HTTPS |
| | | verifiedpermissions.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | verifiedpermissions.us-east-1.amazonaws.com | HTTPS |
| | | verifiedpermissions-fips.us-east-1.api.aws | HTTPS |
| | | verifiedpermissions-fips.us-east-1.amazonaws.com | HTTPS |
| | | | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | verifiedpermissions.us-east-1.api.aws | |
| US West (N. California) | us-west-1 | verifiedpermissions.us-west-1.amazonaws.com | HTTPS |
| | | verifiedpermissions-fips.us-west-1.api.aws | HTTPS |
| | | verifiedpermissions-fips.us-west-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.us-west-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | verifiedpermissions.us-west-2.amazonaws.com | HTTPS |
| | | verifiedpermissions-fips.us-west-2.api.aws | HTTPS |
| | | verifiedpermissions-fips.us-west-2.amazonaws.com | HTTPS |
| | | verifiedpermissions.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | verifiedpermissions.af-south-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | verifiedpermissions.ap-east-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | verifiedpermissions.ap-south-2.amazonaws.com | HTTPS |
| | | verifiedpermissions.ap-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Jakarta) | ap-southeast-3 | verifiedpermissions.ap-southeast-3.amazonaws.com<br><br>verifiedpermissions.ap-southeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | verifiedpermissions.ap-southeast-5.amazonaws.com<br><br>verifiedpermissions.ap-southeast-5.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | verifiedpermissions.ap-southeast-4.amazonaws.com<br><br>verifiedpermissions.ap-southeast-4.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | verifiedpermissions.ap-south-1.amazonaws.com<br><br>verifiedpermissions.ap-south-1.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | verifiedpermissions.ap-southeast-6.amazonaws.com<br><br>verifiedpermissions.ap-southeast-6.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | verifiedpermissions.ap-northeast-3.amazonaws.com<br><br>verifiedpermissions.ap-northeast-3.api.aws | HTTPS<br><br>HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | verifiedpermissions.ap-northeast-2.amazonaws.com<br><br>verifiedpermissions.ap-northeast-2.api.aws | HTTPS<br><br>HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | verifiedpermissions.ap-southeast-1.amazonaws.com<br>verifiedpermissions.ap-southeast-1.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | verifiedpermissions.ap-southeast-2.amazonaws.com<br>verifiedpermissions.ap-southeast-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | verifiedpermissions.ap-east-2.amazonaws.com<br>verifiedpermissions.ap-east-2.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | verifiedpermissions.ap-southeast-7.amazonaws.com<br>verifiedpermissions.ap-southeast-7.api.aws | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | verifiedpermissions.ap-northeast-1.amazonaws.com<br>verifiedpermissions.ap-northeast-1.api.aws | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | verifiedpermissions.ca-central-1.amazonaws.com<br>verifiedpermissions-fips.ca-central-1.api.aws<br>verifiedpermissions-fips.ca-central-1.amazonaws.com<br>verifiedpermissions.ca-central-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Canada West (Calgary) | ca-west-1 | verifiedpermissions.ca-west-1.amazonaws.com | HTTPS |
| | | verifiedpermissions-fips.ca-west-1.api.aws | HTTPS |
| | | verifiedpermissions-fips.ca-west-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.ca-west-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | verifiedpermissions.eu-central-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | verifiedpermissions.eu-west-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | verifiedpermissions.eu-west-2.amazonaws.com | HTTPS |
| | | verifiedpermissions.eu-west-2.api.aws | HTTPS |
| Europe (Milan) | eu-south-1 | verifiedpermissions.eu-south-1.amazonaws.com | HTTPS |
| | | verifiedpermissions.eu-south-1.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | verifiedpermissions.eu-west-3.amazonaws.com | HTTPS |
| | | verifiedpermissions.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | verifiedpermissions.eu-south-2.amazonaws.com | HTTPS |
| | | verifiedpermissions.eu-south-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | verifiedpermissions.eu-north-1.amazonaws.com<br>verifiedpermissions.eu-north-1.api.aws | HTTPS<br>HTTPS |
| Europe (Zurich) | eu-central-2 | verifiedpermissions.eu-central-2.amazonaws.com<br>verifiedpermissions.eu-central-2.api.aws | HTTPS<br>HTTPS |
| Israel (Tel Aviv) | il-central-1 | verifiedpermissions.il-central-1.amazonaws.com<br>verifiedpermissions.il-central-1.api.aws | HTTPS<br>HTTPS |
| Mexico (Central) | mx-central-1 | verifiedpermissions.mx-central-1.amazonaws.com<br>verifiedpermissions.mx-central-1.api.aws | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | verifiedpermissions.me-south-1.amazonaws.com<br>verifiedpermissions.me-south-1.api.aws | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | verifiedpermissions.me-central-1.amazonaws.com<br>verifiedpermissions.me-central-1.api.aws | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | verifiedpermissions.sa-east-1.amazonaws.com<br>verifiedpermissions.sa-east-1.api.aws | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 | verifiedpermissions.us-gov-east-1.amazonaws.com<br>verifiedpermissions-fips.us-gov-east-1.api.aws<br>verifiedpermissions-fips.us-gov-east-1.amazonaws.com<br>verifiedpermissions.us-gov-east-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | verifiedpermissions.us-gov-west-1.amazonaws.com<br>verifiedpermissions-fips.us-gov-west-1.api.aws<br>verifiedpermissions-fips.us-gov-west-1.amazonaws.com<br>verifiedpermissions.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| BatchGetPolicy requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of BatchGetPolicy requests per second per policy store. |
| BatchIsAuthorized requests per second per Region per policy store | Each supported Region: 30 | Yes | The maximum number of BatchIsAuthorized requests per second per policy store. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| BatchIsAuthorizedWithToken requests per second per Region per policy store | Each supported Region: 30 | Yes | The maximum number of BatchIsAuthorizedWithToken requests per second per policy store. |
| CreateIdentitySource requests per second per Region per policy store | Each supported Region: 1 | Yes | The maximum number of CreateIdentitySource requests per second per policy store. |
| CreatePolicy requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of CreatePolicy requests per second per policy store. |
| CreatePolicyStore requests per second per Region per account | Each supported Region: 1 | No | The maximum number of CreatePolicyStore requests per second. |
| CreatePolicyStoreAlias requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of CreatePolicyStoreAlias requests per second per policy store. |
| CreatePolicyTemplate requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of CreatePolicyTemplate requests per second per policy store. |
| DeleteIdentitySource requests per second per Region per policy store | Each supported Region: 1 | Yes | The maximum number of DeleteIdentitySource requests per second per policy store. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| DeletePolicy requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of DeletePolicy requests per second per policy store. |
| DeletePolicyStore requests per second per Region per account | Each supported Region: 1 | No | The maximum number of DeletePolicyStore requests per second. |
| DeletePolicyStoreAlias requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of DeletePolicyStoreAlias requests per second per policy store. |
| DeletePolicyTemplate requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of DeletePolicyTemplate requests per second per policy store. |
| GetIdentitySource requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of GetIdentitySource requests per second per policy store. |
| GetPolicy requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of GetPolicy requests per second per policy store. |
| GetPolicyStore requests per second per Region per account | Each supported Region: 10 | Yes | The maximum number of GetPolicyStore requests per second. |
| GetPolicyStoreAlias requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of GetPolicyStoreAlias requests per second per policy store. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| GetPolicyTemplate requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of GetPolicyTemplate requests per second per policy store. |
| GetSchema requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of GetSchema requests per second per policy store. |
| IsAuthorized requests per second per Region per policy store | Each supported Region: 200 | Yes | The maximum number of IsAuthorized requests per second per policy store. |
| IsAuthorizedWithToken requests per second per Region per policy store | Each supported Region: 200 | Yes | The maximum number of IsAuthorizedWithToken requests per second per policy store. |
| ListIdentitySources requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of ListIdentitySources requests per second per policy store. |
| ListPolicies requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of ListPolicies requests per second per policy store. |
| ListPolicyStoreAliases requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of ListPolicyStoreAliases requests per second per policy store. |
| ListPolicyStores requests per second per Region per account | Each supported Region: 10 | Yes | The maximum number of ListPolicyStores requests per second. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| ListPolicyTemplates requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of ListPolicyTemplates requests per second per policy store. |
| Policy stores per Region per account | Each supported Region: 30,000 | Yes | The maximum number of policy stores. |
| Policy templates per policy store | Each supported Region: 40 | Yes | The maximum number of policy templates in a policy store. |
| PutSchema requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of PutSchema requests per second per policy store. |
| UpdateIdentitySource requests per second per Region per policy store | Each supported Region: 1 | Yes | The maximum number of UpdateIdentitySource requests per second per policy store. |
| UpdatePolicy requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of UpdatePolicy requests per second per policy store. |
| UpdatePolicyStore requests per second per Region per account | Each supported Region: 10 | No | The maximum number of UpdatePolicyStore requests per second. |
| UpdatePolicyTemplate requests per second per Region per policy store | Each supported Region: 10 | Yes | The maximum number of UpdatePolicyTemplate requests per second per policy store. |

# Amazon Virtual Private Cloud endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

The API actions to manage Amazon VPC resources (for example, virtual private clouds, subnets, and gateways) are part of the Amazon EC2 API. For more information, see Amazon VPC actions in the *Amazon EC2 API Reference*.

For the service endpoints for Amazon EC2, see the section called “Amazon EC2”.

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Active VPC peering connections per VPC | Each supported Region: 50 | Yes | The maximum number of active VPC peering connections per VPC. This quota can be increased up to a maximum of 125. |
| Characters per VPC endpoint policy | Each supported Region: 20,480 | No | The maximum number of characters in a VPC endpoint policy, including white space. |
| Egress-only internet gateways per Region | Each supported Region: 5 | Yes | The maximum number of egress-only (outbound- |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | only) internet gateways per Region. This quota is directly tied to the maximum number of VPCs per Region. To increase this quota, increase the number of VPCs per Region. |
| Elastic IP address quota per NAT gateway | Each supported Region: 2 | Yes | The maximum number of Elastic IP addresses that can be associated with a single NAT Gateway of connectivity type public. |
| Gateway VPC endpoints per Region | Each supported Region: 20 | Yes | The maximum number of gateway VPC endpoints per Region. The maximum is 255 gateway endpoints per VPC. |
| IPv4 CIDR blocks per VPC | Each supported Region: 5 | Yes | The maximum number of IPv4 CIDR blocks per VPC. The primary CIDR block and all secondary CIDR blocks count toward this quota. This quota can be increased up to a maximum of 50. |
| IPv6 CIDR blocks per VPC | Each supported Region: 5 | Yes | The maximum number of IPv6 CIDR blocks per VPC. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Inbound or outbound rules per security group | Each supported Region: 60 | Yes | The maximum number of inbound or outbound rules per VPC security group (120 rules in total). This quota is enforced separately for IPv4 and IPv6 rules. A rule that references a security group or prefix list ID counts as one rule each for IPv4 and IPv6. This quota multiplied by the security groups per network interface quota cannot exceed 1000. |
| Interface VPC endpoints per VPC | Each supported Region: 50 | Yes | The maximum number of interface and gatewayLoadBalancer VPC endpoints per VPC. |
| Internet gateways per Region | Each supported Region: 5 | Yes | The maximum number of internet gateways per Region. This quota is directly tied to the maximum number of VPCs per Region. To increase this quota, increase the number of VPCs per Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| NAT gateways per Availability Zone | Each supported Region: 5 | Yes | The maximum number of NAT gateways per Availability Zone. This includes NAT gateways in the pending, active, or deleting state. |
| Network ACLs per VPC | Each supported Region: 200 | Yes | The maximum number of network ACLs per VPC. |
| Network Address Usage | Each supported Region: 64,000 | Yes | The maximum Network Address Usage for a single VPC. |
| Network interfaces per Region | Each supported Region: 5,000 | Yes | The maximum number of network interfaces per Availability Zone in a Region. |
| Outstanding VPC peering connection requests | Each supported Region: 25 | Yes | The maximum number of outstanding VPC peering connection requests that youve requested. |
| Participant accounts per VPC | Each supported Region: 100 | Yes | The maximum number of distinct participant accounts that subnets in a VPC can be shared with. This is a per VPC quota and applies across all the subnets shared in a VPC. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Peered Network Address Usage | Each supported Region: 128,000 | Yes | The maximum Network Address Usage for a VPC and its peers. |
| Private IP address quota per NAT gateway | Each supported Region: 8 | Yes | The maximum number of private IP addresses that can be assigned to a single NAT Gateway of connectivity type private. |
| Regional NAT gateways per VPC | Each supported Region: 5 | Yes | The maximum number of regional NAT gateways per VPC. This includes regional NAT gateways in the pending, active, or deleting state. |
| Route tables per VPC | Each supported Region: 200 | Yes | The maximum number of route tables per VPC. The main route table counts toward this quota. |
| Routes per route table | Each supported Region: 500 | Yes | The maximum number of non-propagated routes per route table. This quota can be increased up to a maximum of 1000; however, network performance might be impacted. This quota is enforced separately for IPv4 and IPv6 routes. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Rules per network ACL | Each supported Region: 20 | Yes | The maximum number of inbound rules or outbound rules per network ACL (a total of 40 rules). This includes both IPv4 and IPv6 rules, and the default deny rules. This quota can be increased up to a maximum of 40; however, network performance might be impacted. |
| Security groups per network interface | Each supported Region: 5 | Yes | The maximum number of security groups per network interface. The maximum is 16. This quota, multiplied by the quota for rules per security group, cannot exceed 1000. |
| Subnets per VPC | Each supported Region: 200 | Yes | The maximum number of subnets per VPC. |
| Subnets that can be shared with an account | Each supported Region: 100 | Yes | The maximum number of subnets that can be shared with an AWS account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| VPC Block Public Access exclusions per account per Region | Each supported Region: 50 | No | The maximum number of VPC Block Public Access exclusions that an account can have in a single Region at a time. |
| VPC associations per security group | Each supported Region: 100 | Yes | The maximum number of VPC associations per security group. |
| VPC endpoints of type resource per VPC | Each supported Region: 500 | Yes | The maximum number of single resource endpoints in a VPC. For additional capacity and limit increases, contact AWS Support. |
| VPC endpoints of type service network per VPC | Each supported Region: 500 | Yes | The maximum number of service network endpoints in a VPC. For additional capacity and limit increases, contact AWS Support. |
| VPC peering connection request expiry hours | Each supported Region: 168 | No | The maximum number of hours after which an unaccepted VPC peering connection request expires. The default value is 168 hours (one week). |
| VPC security groups per Region | Each supported Region: 2,500 | Yes | The maximum number of VPC security groups per Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| VPCs per Region | Each supported Region: 5 | Yes | The maximum number of VPCs per Region. This quota is directly tied to the maximum number of internet gateways per Region. |

> **Note**
>
> The *VPC Block Public Access exclusions per account per Region* quota is adjustable. To request an adjustment, open a service limit increase case using the AWS Support Center Console.

For more information, see the following:

- Amazon VPC quotas
- VPC peering quotas
- Traffic Mirroring quotas
- AWS Transit Gateway quotas
- AWS Network Manager quotas
- Reachability Analyzer quotas
- Network Access Analyzer quotas
- AWS PrivateLink quotas
- AWS Client VPN quotas
- Site-to-Site VPN quotas

## Amazon VPC Lattice endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | vpc-lattice.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | vpc-lattice.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | vpc-lattice.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | vpc-lattice.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | vpc-lattice.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | vpc-lattice.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific | ap-south-2 | vpc-lattice.ap-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| (Hyderabad) | | | | |
| Asia Pacific (Jakarta) | ap-southeast-3 | vpc-lattice.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | vpc-lattice.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | vpc-lattice.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | vpc-lattice.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | vpc-lattice.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | vpc-lattice.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | vpc-lattice.ap-southeast-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | vpc-lattice.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | vpc-lattice.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | vpc-lattice.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | vpc-lattice.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | vpc-lattice.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | vpc-lattice.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | vpc-lattice.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | vpc-lattice.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | vpc-lattice.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | vpc-lattice.eu-south-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | vpc-lattice.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | vpc-lattice.eu-central-2.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | vpc-lattice.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | vpc-lattice.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | vpc-lattice.me-central-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | vpc-lattice.sa-east-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Auth policy size | Each supported Region: 10 Kilobytes | No | The maximum size of a JSON file in an Auth policy. |
| Child Resource Configurations per Group Resource Configuration | Each supported Region: 60 | Yes | The maximum number of child resource configura |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | tions in a group resource configuration. For additional capacity and limit increases, contact AWS Support. |
| Domain Verifications per AWS Region | Each supported Region: 5 | Yes | The maximum number of domain verifications that can be created per account. For additional capacity and limit increases, contact AWS Support. |
| Listeners per service | Each supported Region: 2 | Yes | The maximum number of listeners that you can create for a service. For additional capacity and limit increases, contact AWS Support. |
| Resource Configurations per service network | Each supported Region: 500 | Yes | The maximum number of resource configurations associated with a service network. For additional capacity and limit increases, contact AWS Support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Resource configurations per AWS Region | Each supported Region: 2,000 | Yes | The maximum number of resource configurations an AWS account can have per AWS Region. For additional capacity and limit increases, contact AWS Support. |
| Resource gateways per VPC | Each supported Region: 500 | Yes | The maximum number of resource gateways in a VPC. For additional capacity and limit increases, contact AWS Support. |
| Rules per listener | Each supported Region: 10 | Yes | The maximum number of rules that you can define for your service listener. For additional capacity and limit increases, contact AWS Support. |
| Security groups per association | Each supported Region: 5 | No | The maximum number of security groups that you can add to an association between a VPC and a service network. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Service associations per service network | Each supported Region: 500 | Yes | The maximum number of services that you can associate with a single service network. For additional capacity and limit increases, contact AWS Support. |
| Service networks per region | Each supported Region: 50 | Yes | The maximum number of service networks per region. For additional capacity and limit increases, contact AWS Support. |
| Services per region | Each supported Region: 2,000 | Yes | The maximum number of services per region. For additional capacity and limit increases, contact AWS Support. |
| Target groups per region | Each supported Region: 500 | Yes | The maximum number of target groups per region. For additional capacity and limit increases, contact AWS Support. |
| Target groups per service | Each supported Region: 10 | Yes | The maximum number of target groups that you can associate with a service. For additional capacity and limit increases, contact AWS Support. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Targets per target group | Each supported Region: 1,000 | Yes | The maximum number of targets that you can associate with a single target group. For additional capacity and limit increases, contact AWS Support. |
| VPC associations per service network | Each supported Region: 500 | Yes | The maximum number of VPCs that you can associate with a single service network. For additional capacity and limit increases, contact AWS Support. |
| VPC endpoints of type service network per service network | Each supported Region: 200 | Yes | The maximum number of service network endpoints associated with a service network. For additional capacity and limit increases, contact AWS Support. |

## AWS WAF endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

> **Note**
>
> This page provides information related the latest version of AWS WAF, released in November 2019. The names of the entities that you use to access AWS WAF, like endpoints and namespaces, all have the versioning information added, like V2 or v2, to distinguish from the prior version.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | wafv2.us-east-2.amazonaws.com | HTTPS |
| | | wafv2-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | wafv2.us-east-1.amazonaws.com | HTTPS |
| | | wafv2-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | wafv2.us-west-1.amazonaws.com | HTTPS |
| | | wafv2-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | wafv2.us-west-2.amazonaws.com | HTTPS |
| | | wafv2-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | wafv2.af-south-1.amazonaws.com | HTTPS |
| | | wafv2-fips.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific | ap-east-1 | wafv2.ap-east-1.amazonaws.com | HTTPS |
| | | wafv2-fips.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| (Hong Kong) | | | |
| Asia Pacific (Hyderabad) | ap-south-2 | wafv2.ap-south-2.amazonaws.com<br>wafv2-fips.ap-south-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | wafv2.ap-southeast-3.amazonaws.com<br>wafv2-fips.ap-southeast-3.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | wafv2.ap-southeast-5.amazonaws.com<br>wafv2-fips.ap-southeast-5.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | wafv2.ap-southeast-4.amazonaws.com<br>wafv2-fips.ap-southeast-4.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | wafv2.ap-south-1.amazonaws.com<br>wafv2-fips.ap-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (New Zealand) | ap-southeast-6 | wafv2.ap-southeast-6.amazonaws.com<br>wafv2-fips.ap-southeast-6.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | wafv2.ap-northeast-3.amazonaws.com<br>wafv2-fips.ap-northeast-3.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Seoul) | ap-northeast-2 | wafv2.ap-northeast-2.amazonaws.com<br>wafv2-fips.ap-northeast-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | wafv2.ap-southeast-1.amazonaws.com<br>wafv2-fips.ap-southeast-1.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | wafv2.ap-southeast-2.amazonaws.com<br>wafv2-fips.ap-southeast-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Taipei) | ap-east-2 | wafv2.ap-east-2.amazonaws.com<br>wafv2-fips.ap-east-2.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Thailand) | ap-southeast-7 | wafv2.ap-southeast-7.amazonaws.com<br>wafv2-fips.ap-southeast-7.amazonaws.com | HTTPS<br>HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | wafv2.ap-northeast-1.amazonaws.com<br>wafv2-fips.ap-northeast-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada (Central) | ca-central-1 | wafv2.ca-central-1.amazonaws.com<br>wafv2-fips.ca-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Canada West (Calgary) | ca-west-1 | wafv2.ca-west-1.amazonaws.com<br>wafv2-fips.ca-west-1.amazonaws.com | HTTPS<br>HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | wafv2.eu-central-1.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | wafv2.eu-west-1.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | wafv2.eu-west-2.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | wafv2.eu-south-1.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | wafv2.eu-west-3.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | wafv2.eu-south-2.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-south-2.amazonaws.com | HTTPS |
| Europe (Stockholm) | eu-north-1 | wafv2.eu-north-1.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | wafv2.eu-central-2.amazonaws.com | HTTPS |
| | | wafv2-fips.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | wafv2.il-central-1.amazonaws.com | HTTPS |
| | | wafv2-fips.il-central-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Mexico (Central) | mx-central-1 | wafv2.mx-central-1.amazonaws.com<br>wafv2-fips.mx-central-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (Bahrain) | me-south-1 | wafv2.me-south-1.amazonaws.com<br>wafv2-fips.me-south-1.amazonaws.com | HTTPS<br>HTTPS |
| Middle East (UAE) | me-central-1 | wafv2.me-central-1.amazonaws.com<br>wafv2-fips.me-central-1.amazonaws.com | HTTPS<br>HTTPS |
| South America (São Paulo) | sa-east-1 | wafv2.sa-east-1.amazonaws.com<br>wafv2-fips.sa-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | wafv2.us-gov-east-1.amazonaws.com<br>wafv2-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | wafv2.us-gov-west-1.amazonaws.com<br>wafv2-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum IP sets per account in WAF for CloudFront | Each supported Region: 100 | Yes | The maximum number of IP sets you can create |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | in your account for CloudFront. |
| Maximum IP sets per account in WAF for regional | Each supported Region: 100 | Yes | The maximum number of IP sets you can create in your account for regional. |
| Maximum combined size in kilobytes of all response body content for a single rule group or a single web ACL for CloudFront | Each supported Region: 50 | No | The maximum combined size in kilobytes of all response body content for a single rule group or a single web ACL for CloudFront. |
| Maximum combined size in kilobytes of all response body content for a single rule group or a single web ACL for regional | Each supported Region: 50 | No | The maximum combined size in kilobytes of all response body content for a single rule group or a single web AC for regional. |
| Maximum number of ALB associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of ALB associations per web ACL for regional. |
| Maximum number of AWS Amplify associations per web ACL in WAF for CloudFront | Each supported Region: 100 | Yes | The maximum number of AWS Amplify associations per web ACL for CloudFront. |
| Maximum number of AWS App Runner service associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of AWS App Runner service associations per web ACL for regional. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of AWS AppSync GraphQL API associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of AWS AppSync GraphQL API associations per web ACL for regional. |
| Maximum number of AWS Verified Access instance associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of AWS Verified Access instance associations per web ACL for regional. |
| Maximum number of Amazon API Gateway REST API associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of Amazon API Gateway REST API associations per web ACL for regional. |
| Maximum number of Amazon Bedrock AgentCore Gateway associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of Amazon Bedrock AgentCore Gateway associations per web ACL for regional. |
| Maximum number of Amazon Cognito user pool associations per web ACL in WAF for regional | Each supported Region: 100 | Yes | The maximum number of Amazon Cognito user pool associations per web ACL for regional. |
| Maximum number of IP addresses in an IP set in WAF for CloudFront | Each supported Region: 10,000 | No | The maximum number of IP addresses allowed in an IP set for CloudFront. |
| Maximum number of IP addresses in an IP set in WAF for regional | Each supported Region: 10,000 | No | The maximum number of IP addresses allowed in an IP set for regional. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of bytes in a string match (byte match) string in WAF for CloudFront | Each supported Region: 200 | No | The maximum number of bytes you can create in a string match (byte match) string for CloudFront. |
| Maximum number of bytes in a string match (byte match) string in WAF for regional | Each supported Region: 200 | No | The maximum number of bytes you can create in a string match (byte match) string for regional. |
| Maximum number of characters allowed in a regex pattern per account in WAF for Cloudfront | Each supported Region: 200 | No | The maximum number of characters allowed in a regex pattern in your account for CloudFront. |
| Maximum number of characters allowed in a regex pattern per account in WAF for regional | Each supported Region: 200 | No | The maximum number of characters allowed in a regex pattern in your account for regional. |
| Maximum number of custom headers for a single custom request definition for CloudFront | Each supported Region: 10 | No | The maximum number of custom headers for a single custom request definition for CloudFront. |
| Maximum number of custom headers for a single custom request definition for regional | Each supported Region: 10 | No | The maximum number of custom headers for a single custom request definition for regional. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of custom headers for a single custom response definition for CloudFront | Each supported Region: 10 | No | The maximum number of custom headers for a single custom response definition for CloudFront. |
| Maximum number of custom headers for a single custom response definition for regional | Each supported Region: 10 | No | The maximum number of custom headers for a single custom response definition for regional. |
| Maximum number of custom request headers per web ACL or rule group for CloudFront | Each supported Region: 100 | Yes | The maximum number of custom request headers per web ACL or rule group for CloudFront. |
| Maximum number of custom request headers per web ACL or rule group for regional | Each supported Region: 100 | Yes | The maximum number of custom request headers per web ACL or rule group for regional. |
| Maximum number of custom response bodies per web ACL or rule group for CloudFront | Each supported Region: 50 | Yes | The maximum number of custom response bodies per web ACL or rule group for CloudFront. |
| Maximum number of custom response bodies per web ACL or rule group for regional | Each supported Region: 50 | Yes | The maximum number of custom response bodies per web ACL or rule group for regional. |
| Maximum number of custom response headers per web ACL or rule group for CloudFront | Each supported Region: 100 | Yes | The maximum number of custom response headers per web ACL or rule group for CloudFront. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of custom response headers per web ACL or rule group for regional | Each supported Region: 100 | Yes | The maximum number of custom response headers per web ACL or rule group for regional. |
| Maximum number of patterns in a regex pattern set per account in WAF for Cloudfront | Each supported Region: 10 | No | The maximum number of patterns in a regex pattern set you can create in your account for CloudFront. |
| Maximum number of patterns in a regex pattern set per account in WAF for regional | Each supported Region: 10 | No | The maximum number of patterns in a regex pattern set you can create in your account for regional. |
| Maximum number of rate-based statements per rule group in WAF for Cloudfront | Each supported Region: 4 | No | The maximum number of rate-based statements allowed in a rule group for CloudFront. |
| Maximum number of rate-based statements per rule group in WAF for regional | Each supported Region: 4 | No | The maximum number of rate-based statements allowed in a rule group for regional. |
| Maximum number of rate-based statements per web ACL in WAF for Cloudfront | Each supported Region: 10 | No | The maximum number of rate-based statements allowed in a web ACL for CloudFront. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum number of rate-based statements per web ACL in WAF for regional | Each supported Region: 10 | No | The maximum number of rate-based statements allowed in a web ACL for regional. |
| Maximum number of referenced statements per rule group or web ACL in WAF for Cloudfront | Each supported Region: 50 | No | The maximum number of referenced statements allowed within a rule group or web ACL for CloudFront. |
| Maximum number of referenced statements per rule group or web ACL in WAF for regional | Each supported Region: 50 | No | The maximum number of referenced statements allowed within a rule group or web ACL for regional. |
| Maximum number of requests per second per web ACL for CloudFront | Each supported Region: 100,000 | Yes | The maximum number of requests per second per web ACL for CloudFront. |
| Maximum number of requests per second per web ACL for regional | Each supported Region: 100,000 | Yes | The maximum number of requests per second per web ACL for regional. |
| Maximum number of text transformations per rule statement for CloudFront | Each supported Region: 10 | No | The maximum number of text transformations per rule statement for CloudFront. |
| Maximum number of text transformations per rule statement for regional | Each supported Region: 10 | No | The maximum number of text transformations per rule statement for regional. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum number of unique IP addresses that can be blocked per rate-based rule for CloudFront | Each supported Region: 10,000 | No | The maximum number of unique IP addresses that can be blocked per rate-based rule for CloudFront. |
| Maximum number of unique IP addresses that can be blocked per rate-based rule for regional | Each supported Region: 10,000 | No | The maximum number of unique IP addresses that can be blocked per rate-based rule for regional. |
| Maximum number of web ACL capacity units in a rule group in WAF for CloudFront | Each supported Region: 5,000 | No | The maximum number of web ACL capacity units allowed in a rule group for CloudFront. |
| Maximum number of web ACL capacity units in a rule group in WAF for regional | Each supported Region: 5,000 | No | The maximum number of web ACL capacity units allowed in a rule group for regional. |
| Maximum number of web ACL capacity units in a web ACL in WAF for CloudFront | Each supported Region: 5,000 | No | The maximum number of web ACL capacity units allowed in a web ACL for CloudFront. |
| Maximum number of web ACL capacity units in a web ACL in WAF for regional | Each supported Region: 5,000 | No | The maximum number of web ACL capacity units allowed in a web ACL for regional. |
| Maximum regex pattern sets per account in WAF for CloudFront | Each supported Region: 10 | Yes | The maximum number of regex pattern sets you can create in your account for CloudFront. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum regex pattern sets per account in WAF for regional | Each supported Region: 10 | Yes | The maximum number of regex pattern sets you can create in your account for regional. |
| Maximum rule groups per account in WAF for CloudFront | Each supported Region: 100 | Yes | The maximum number of rule groups you can create in your account for CloudFront. |
| Maximum rule groups per account in WAF for regional | Each supported Region: 100 | Yes | The maximum number of rule groups you can create in your account for regional. |
| Maximum size in kilobytes of a web request body that can be inspected for Application Load Balancer and AWS AppSync protections | Each supported Region: 8 | No | The maximum size in kilobytes of a web request body that can be inspected for Application Load Balancer and AWS AppSync protections. |
| Maximum size in kilobytes of a web request body that can be inspected for CloudFront, API Gateway, Amazon Cognito, App Runner, and Verified Access protections | Each supported Region: 64 | No | The maximum size in kilobytes of a web request body that can be inspected for CloudFront, API Gateway, Amazon Cognito, App Runner, and Verified Access protections. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Maximum size in kilobytes of the custom response body content for a single custom response definition for CloudFront | Each supported Region: 4 | No | The maximum size in kilobytes of the custom response body content for a single custom response definition for CloudFront. |
| Maximum size in kilobytes of the custom response body content for a single custom response definition for regional | Each supported Region: 4 | No | The maximum size in kilobytes of the custom response body content for a single custom response definition for regional. |
| Maximum web ACLs per account in WAF for CloudFront | Each supported Region: 100 | Yes | The maximum number of web ACLs you can create in your account for CloudFront. |
| Maximum web ACLs per account in WAF for regional | Each supported Region: 100 | Yes | The maximum number of web ACLs you can create in your account for regional. |
| Minimum request rate that can be defined for a rate-based rule for CloudFront | Each supported Region: 10 | No | The minimum request rate that can be defined for a rate-based rule for CloudFront. |
| Minimum request rate that can be defined for a rate-based rule for regional | Each supported Region: 10 | No | The minimum request rate that can be defined for a rate-based rule for regional. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of token domains per web ACL for CloudFront | Each supported Region: 10 | Yes | The number of token domains per web ACL for CloudFront. |
| Number of token domains per web ACL for regional | Each supported Region: 10 | Yes | The number of token domains per web ACL for regional. |

For more information, see AWS WAF quotas in the *AWS WAF Developer Guide*.

## AWS WAF Classic endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

**Note**

This page provides information related to **AWS WAF Classic**. If you created AWS WAF resources, like rules and web ACLs, in AWS WAF prior to November 2019, and you have not migrated your web ACLs over yet, you must use AWS WAF Classic to access those resources. Otherwise, do not use this version.
**For information related to the latest version of AWS WAF**, see AWS WAF endpoints and quotas.

# Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | waf.amazonaws.com | HTTPS | |
| | | waf-fips.amazonaws.com | HTTPS | |

AWS WAF Classic for Application Load Balancers and API Gateway APIs has the following endpoints:

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | waf-regional.us-east-2.amazonaws.com | HTTPS | |
| | | waf-regional-fips.us-east-2.amazonaws.com | HTTPS | |
| US East (N. Virginia) | us-east-1 | waf-regional.us-east-1.amazonaws.com | HTTPS | |
| | | waf-regional-fips.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | waf-regional.us-west-1.amazonaws.com | HTTPS | |
| | | waf-regional-fips.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | waf-regional.us-west-2.amazonaws.com | HTTPS | |
| | | waf-regional-fips.us-west-2.amazonaws.com | HTTPS | |
| Africa (Cape Town) | af-south-1 | waf-regional.af-south-1.amazonaws.com | HTTPS | |
| | | waf-regional-fips.af-south-1.amazonaws.com | HTTPS | |
| Asia Pacific | ap-east-1 | waf-regional.ap-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| (Hong Kong) | | waf-regional-fips.ap-east-1.amazonaws.com | HTTPS |
| Asia Pacific (Hyderabad) | ap-south-2 | waf-regional.ap-south-2.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-south-2.amazonaws.com | HTTPS |
| Asia Pacific (Jakarta) | ap-southeast-3 | waf-regional.ap-southeast-3.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-southeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Melbourne) | ap-southeast-4 | waf-regional.ap-southeast-4.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-southeast-4.amazonaws.com | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | waf-regional.ap-south-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Osaka) | ap-northeast-3 | waf-regional.ap-northeast-3.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-northeast-3.amazonaws.com | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | waf-regional.ap-northeast-2.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-northeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | waf-regional.ap-southeast-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-southeast-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Asia Pacific (Sydney) | ap-southeast-2 | waf-regional.ap-southeast-2.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-southeast-2.amazonaws.com | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | waf-regional.ap-northeast-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.ap-northeast-1.amazonaws.com | HTTPS |
| Canada (Central) | ca-central-1 | waf-regional.ca-central-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.ca-central-1.amazonaws.com | HTTPS |
| Europe (Frankfurt) | eu-central-1 | waf-regional.eu-central-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-central-1.amazonaws.com | HTTPS |
| Europe (Ireland) | eu-west-1 | waf-regional.eu-west-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-west-1.amazonaws.com | HTTPS |
| Europe (London) | eu-west-2 | waf-regional.eu-west-2.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-west-2.amazonaws.com | HTTPS |
| Europe (Milan) | eu-south-1 | waf-regional.eu-south-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-south-1.amazonaws.com | HTTPS |
| Europe (Paris) | eu-west-3 | waf-regional.eu-west-3.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-west-3.amazonaws.com | HTTPS |
| Europe (Spain) | eu-south-2 | waf-regional.eu-south-2.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-south-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Stockholm) | eu-north-1 | waf-regional.eu-north-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-north-1.amazonaws.com | HTTPS |
| Europe (Zurich) | eu-central-2 | waf-regional.eu-central-2.amazonaws.com | HTTPS |
| | | waf-regional-fips.eu-central-2.amazonaws.com | HTTPS |
| Israel (Tel Aviv) | il-central-1 | waf-regional.il-central-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.il-central-1.amazonaws.com | HTTPS |
| Middle East (Bahrain) | me-south-1 | waf-regional.me-south-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.me-south-1.amazonaws.com | HTTPS |
| Middle East (UAE) | me-central-1 | waf-regional.me-central-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.me-central-1.amazonaws.com | HTTPS |
| South America (São Paulo) | sa-east-1 | waf-regional.sa-east-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | waf-regional.us-gov-east-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.us-gov-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | waf-regional.us-gov-west-1.amazonaws.com | HTTPS |
| | | waf-regional-fips.us-gov-west-1.amazonaws.com | HTTPS |

# Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Conditions per rule | Each supported Region: 10 | No | The maximum number of conditions you can add to a rule. |
| Filters per SQL injection match condition | Each supported Region: 10 | No | The maximum number of filters you can add to a SQL injection match condition. |
| Filters per cross-site scripting match condition | Each supported Region: 10 | No | The maximum number of filters you can add to a cross-site scripting match condition. |
| Filters per size constraint condition | Each supported Region: 10 | No | The maximum number of filters you can add to a size constraint condition. |
| Filters per string match condition | Each supported Region: 10 | No | The maximum number of filters you can add to a string match condition. |
| GeoMatchSets | Each supported Region: 50 | No | The maximum number of GeoMatchSets you can add to your account. |
| HTTP header name length | Each supported Region: 40 | No | The length, in bytes, that you want AWS WAF to inspect the HTTP header for in a size constraint condition. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| IP address ranges per IP set match condition | Each supported Region: 10,000 | No | The maximum number of IP address ranges (in CIDR notation) you can add to an IP Set match condition. |
| IP addresses blocked per rate-based rule | Each supported Region: 10,000 | No | The maximum number of IP addresses blocked per rate-based rule. |
| Locations per GeoMatchSet | Each supported Region: 50 | No | The maximum number of locations you can add to a GeoMatchSet. |
| Logging destination configurations per web ACL | Each supported Region: 1 | No | The maximum number of logging destination configurations you can add to a web ACL. |
| Pattern sets per regex match condition | Each supported Region: 1 | No | The maximum number of pattern sets you can add to regex match condition . |
| Patterns per pattern set | Each supported Region: 10 | No | The maximum number of patterns you can add to a pattern set in a regex match condition. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rate of requests | Each supported Region: 10,000 | Yes | The maximum number of requests per second on an Application Load Balancer. This limit applies only to AWS WAF on an Application Load Balancer. This limit can be increased using the AWS console. |
| Rate-based rule rate | Each supported Region: 2,000 | No | The maximum number of requests from a single IP address allowed in a five-minute period. |
| Rate-based rules | Each supported Region: 5 | Yes | The maximum number of rate-based rules you can create in your account. |
| Regex pattern length | Each supported Region: 70 | No | The number of characters that you want AWS WAF to search for in the pattern in a regex match condition. |
| Regex pattern sets | Each supported Region: 5 | No | The maximum number of regex pattern sets you can add to your account. |
| Rules | Each supported Region: 100 | Yes | The maximum number of rules you can create in your account. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Rules per web ACL | Each supported Region: 10 | No | The maximum number of rules you can add to a web ACL. |
| Search length | Each supported Region: 50 | No | The length, in bytes, that you want AWS WAF to watch for in a size constraint condition. |
| Web ACLs | Each supported Region: 50 | Yes | The maximum number of Web ACLs you can create in your account. |

For more information, see AWS WAF Classic quotas in the *AWS WAF Developer Guide*.

# AWS Well-Architected Tool endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (Ohio) | us-east-2 | wellarchitected.us-east-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| US East (N. Virginia) | us-east-1 | wellarchitected.us-east-1.amazonaws.com | HTTPS | |
| US West (N. California) | us-west-1 | wellarchitected.us-west-1.amazonaws.com | HTTPS | |
| US West (Oregon) | us-west-2 | wellarchitected.us-west-2.amazonaws.com | HTTPS | |
| Asia Pacific (Hong Kong) | ap-east-1 | wellarchitected.ap-east-1.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | wellarchitected.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | wellarchitected.ap-northeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | wellarchitected.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | wellarchitected.ap-southeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Tokyo) | ap-northeast-1 | wellarchitected.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | wellarchitected.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | wellarchitected.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | wellarchitected.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | wellarchitected.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | wellarchitected.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | wellarchitected.eu-north-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | wellarchitected.me-south-1.amazonaws.com | HTTPS | |
| South America (São Paulo) | sa-east-1 | wellarchitected.sa-east-1.amazonaws.com | HTTPS | |
| AWS GovCloud (US-East) | us-gov-east-1 | wellarchitected.us-gov-east-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | wellarchitected.us-gov-west-1.amazonaws.com | HTTPS | |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Applications per profile | Each supported Region: 0 | No | The maximum number of applications that can be associated with a Well-Architected Agent profile. |
| Architecture recommendation generations per day | Each supported Region: 5 | No | Maximum number of on-demand architecture recommendation generations per day per account. |
| Choices per question | Each supported Region: 15 | No | The maximum number of choices that can be created for a question. |
| Lens size | Each supported Region: 500 Kilobytes | No | The maximum lens size, in KB. |
| Lenses per account per Region | Each supported Region: 15 | No | The maximum number of lenses that can be |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| | | | created per account in a Region. |
| Lenses per review template | Each supported Region: 20 | No | The maximum number of lenses that can be associated with a review template. |
| Lenses per workload | Each supported Region: 20 | No | The maximum number of lenses that can be associated with a workload. |
| Milestones per workload | Each supported Region: 100 | No | The maximum number of milestones that can be created for a workload. |
| Pillars per lens | Each supported Region: 10 | No | The maximum number of pillars that can be created for a lens. |
| Profiles per account | Each supported Region: 0 | No | The maximum number of Well-Architected Agent profiles per AWS account. |
| Questions per pillar | Each supported Region: 20 | No | The maximum number of questions that can be created for a pillar. |
| Review templates per account per Region | Each supported Region: 500 | No | The maximum number of review templates that can be created per account in a Region. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Shares per lens | Each supported Region: 300 | No | The maximum number of shares that can be created for a lens. |
| Shares per review template | Each supported Region: 20 | No | The maximum number of shares that can be created for a review template. |
| Shares per workload | Each supported Region: 20 | No | The maximum number of shares that can be created for a workload. |
| Versions per lens | Each supported Region: 100 | No | The maximum number of versions that can be created for a lens. |
| Workloads per account per Region | Each supported Region: 1,000 | No | The maximum number of workloads that can be created per account in a Region. |

# AWS Wickr endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

# Service endpoints

## AWS Wickr Production API

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | admin.wickr.us-east-1.amazonaws.com | HTTPS |
| | | admin.wickr.us-east-1.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | admin.wickr.ap-southeast-1.amazonaws.com | HTTPS |
| | | admin.wickr.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | admin.wickr.ap-southeast-2.amazonaws.com | HTTPS |
| | | admin.wickr.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | admin.wickr.ap-northeast-1.amazonaws.com | HTTPS |
| | | admin.wickr.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | admin.wickr.ca-central-1.amazonaws.com | HTTPS |
| | | admin.wickr.ca-central-1.api.aws | HTTPS |
| Europe (Frankfurt) | eu-central-1 | admin.wickr.eu-central-1.amazonaws.com | HTTPS |
| | | admin.wickr.eu-central-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | admin.wickr.eu-west-2.amazonaws.com | HTTPS |
| | | admin.wickr.eu-west-2.api.aws | HTTPS |
| Europe (Stockholm) | eu-north-1 | admin.wickr.eu-north-1.amazonaws.com | HTTPS |
| | | admin.wickr.eu-north-1.api.aws | HTTPS |
| Europe (Zurich) | eu-central-2 | admin.wickr.eu-central-2.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | admin.wickr.eu-central-2.api.aws | HTTPS |

### AWS Wickr GovCloud API

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | admin.wickr.us-gov-west-1.amazonaws.com | HTTPS |
| | | admin.wickr.us-gov-west-1.api.aws | HTTPS |

# Amazon WorkMail endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Service | Endpoint |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | Amazon WorkMail SDK | https://workmail.us-east-1.amazonaws.com |
| US East (N. Virginia) | us-east-1 | Amazon WorkMail Message Flow SDK | https://workmailmessageflow.us-east-1.amazonaws.com |

| Region Name | Region | Service | Endpoint |
|---|---|---|---|
| US West (Oregon) | us-west-2 | Amazon WorkMail SDK | https://workmail.us-west-2.amazonaws.com |
| US West (Oregon) | us-west-2 | Amazon WorkMail Message Flow SDK | https://workmailmessageflow.us-west-2.amazonaws.com |
| Europe (Ireland) | eu-west-1 | Amazon WorkMail SDK | https://workmail.eu-west-1.amazonaws.com |
| Europe (Ireland) | eu-west-1 | Amazon WorkMail Message Flow SDK | https://workmailmessageflow.eu-west-1.amazonaws.com |

## Email protocols endpoints

| Region Name | Region | Service | Endpoint |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | Autodiscover | https://autodiscover-service.mail.us-east-1.awsapps.com/autodiscover/autodiscover.xml |
| US East (N. Virginia) | us-east-1 | Exchange web services (EWS) | https://ews.mail.us-east-1.awsapps.com/EWS/Exchange.asmx |
| US East (N. Virginia) | us-east-1 | Exchange ActiveSync (EAS) | https://mobile.mail.us-east-1.awsapps.com/Microsoft-Server-ActiveSync |
| US East (N. Virginia) | us-east-1 | IMAP over SSL/TLS (IMAPS) | imap.mail.us-east-1.awsapps.com:993 |
| US East (N. Virginia) | us-east-1 | SMTP over SSL/TLS (SMTPS) | smtp.mail.us-east-1.awsapps.com:465 |
| US West (Oregon) | us-west-2 | Autodiscover | https://autodiscover-service.mail.us-west-2.awsapps.com/autodiscover/autodiscover.xml |

| Region Name | Region | Service | Endpoint |
|---|---|---|---|
| US West (Oregon) | us-west-2 | Exchange web services (EWS) | https://ews.mail.us-west-2.awsapps.com/EWS/Exchange.asmx |
| US West (Oregon) | us-west-2 | Exchange ActiveSync (EAS) | https://mobile.mail.us-west-2.awsapps.com/Microsoft-Server-ActiveSync |
| US West (Oregon) | us-west-2 | IMAP over SSL/TLS (IMAPS) | imap.mail.us-west-2.awsapps.com:993 |
| US West (Oregon) | us-west-2 | SMTP over SSL/TLS (SMTPS) | smtp.mail.us-west-2.awsapps.com:465 |
| Europe (Ireland) | eu-west-1 | Autodiscover | https://autodiscover-service.mail.eu-west-1.awsapps.com/autodiscover/autodiscover.xml |
| Europe (Ireland) | eu-west-1 | Exchange web services (EWS) | https://ews.mail.eu-west-1.awsapps.com/EWS/Exchange.asmx |
| Europe (Ireland) | eu-west-1 | Exchange ActiveSync (EAS) | https://mobile.mail.eu-west-1.awsapps.com/Microsoft-Server-ActiveSync |
| Europe (Ireland) | eu-west-1 | IMAP over SSL/TLS (IMAPS) | imap.mail.eu-west-1.awsapps.com:993 |
| Europe (Ireland) | eu-west-1 | SMTP over SSL/TLS (SMTPS) | smtp.mail.eu-west-1.awsapps.com:465 |

### Service quotas

For more information, see Amazon WorkMail Quotas.

## Amazon WorkSpaces endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

**Note**

The AWS Regions in the following table apply to WorkSpaces personal. For the AWS Regions that apply to WorkSpaces pools, see AWS Regions for WorkSpaces pools in the *Amazon WorkSpaces Administration Guide*. The endpoints are the same for both versions of the WorkSpaces service.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | workspaces.us-east-2.amazonaws.com | HTTPS |
| | | workspaces-fips.us-east-2.api.aws | HTTPS |
| | | workspaces-fips.us-east-2.amazonaws.com | HTTPS |
| | | workspaces.us-east-2.api.aws | HTTPS |
| US East (N. Virginia) | us-east-1 | workspaces.us-east-1.amazonaws.com | HTTPS |
| | | workspaces-fips.us-east-1.api.aws | HTTPS |
| | | workspaces-fips.us-east-1.amazonaws.com | HTTPS |
| | | workspaces.us-east-1.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | workspaces.us-west-2.amazonaws.com | HTTPS |
| | | workspaces-fips.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | workspaces-fips.us-west-2.amazonaws.com | HTTPS |
| | | workspaces.us-west-2.api.aws | HTTPS |
| Africa (Cape Town) | af-south-1 | workspaces.af-south-1.amazonaws.com | HTTPS |
| | | workspaces.af-south-1.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | workspaces.ap-southeast-5.amazonaws.com | HTTPS |
| | | workspaces.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | workspaces.ap-south-1.amazonaws.com | HTTPS |
| | | workspaces.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | workspaces.ap-northeast-2.amazonaws.com | HTTPS |
| | | workspaces.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | workspaces.ap-southeast-1.amazonaws.com | HTTPS |
| | | workspaces.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | workspaces.ap-southeast-2.amazonaws.com | HTTPS |
| | | workspaces.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | workspaces.ap-northeast-1.amazonaws.com | HTTPS |
| | | workspaces.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | workspaces.ca-central-1.amazonaws.com | HTTPS |
| | | workspaces.ca-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | workspaces.eu-central-1.amazonaws.com | HTTPS |
| | | workspaces.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | workspaces.eu-west-1.amazonaws.com | HTTPS |
| | | workspaces.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | workspaces.eu-west-2.amazonaws.com | HTTPS |
| | | workspaces.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | workspaces.eu-west-3.amazonaws.com | HTTPS |
| | | workspaces.eu-west-3.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | workspaces.il-central-1.amazonaws.com | HTTPS |
| | | workspaces.il-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | workspaces.sa-east-1.amazonaws.com | HTTPS |
| | | workspaces.sa-east-1.api.aws | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | workspaces.us-gov-east-1.amazonaws.com | HTTPS |
| | | workspaces-fips.us-gov-east-1.api.aws | HTTPS |
| | | workspaces-fips.us-gov-east-1.amazonaws.com | HTTPS |
| | | workspaces.us-gov-east-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| AWS GovCloud (US-West) | us-gov-west-1 | workspaces.us-gov-west-1.amazonaws.com<br>workspaces-fips.us-gov-west-1.api.aws<br>workspaces-fips.us-gov-west-1.amazonaws.com<br>workspaces.us-gov-west-1.api.aws | HTTPS<br>HTTPS<br>HTTPS<br>HTTPS | |

## Service quotas

| Resource | Default | Description | Adjustable |
|---|---|---|---|
| WorkSpaces | 1 | The maximum number of WorkSpaces in this account in the current Region. | Yes |
| Graphics WorkSpaces | 0 | The maximum number of Graphics WorkSpaces in this account in the current Region.<br>**Note**<br>Graphics bundle is no longer supported after November 30, 2023. We | Yes |

| Resource | Default | Description | Adjustable |
| --- | --- | --- | --- |
| | | recommend migrating your WorkSpaces to Graphics.g4dn bundle. For more information, see Migrate a WorkSpace in the *WorkSpaces Administration Guide*. | |
| GraphicsPro WorkSpaces | 0 | The maximum number of GraphicsPro WorkSpaces in this account in the current Region. | Yes |
| Images | 40 | The maximum number of images in this account in the current Region. | Yes |
| Bundles | 50 | The maximum number of bundles in this account in the current Region. This quota applies only to custom bundles, not to public bundles. | No |

| Resource | Default | Description | Adjustable |
| --- | --- | --- | --- |
| Connection aliases | 20 | The maximum number of connection aliases in this account in the current Region. | No |
| Directories | 50 | The maximum number of directories that can be registered for use with Amazon WorkSpaces in this account in the current Region. | No |
| IP access control groups | 100 | The maximum number of IP access control groups in this account in the current Region. | No |
| Rules per IP access control group | 10 | The maximum number of rules per IP access control group in this account in the current Region. | No |
| IP access control groups per directory | 25 | The maximum number of IP access control groups per directory in this account in the current Region. | No |

| Resource | Default | Description | Adjustable |
| --- | --- | --- | --- |
| WorkSpaces Pools | 10 | The maximum number of WorkSpaces Pools in this account in the current Region. | Yes |
| General Purpose Value streaming instances for WorkSpaces Pools | 10 | The maximum number of General Purpose Value streaming instances that can be used for WorkSpaces Pools in this account in the current Region. | Yes |
| General Purpose Standard streaming instances for WorkSpaces Pools | 10 | The maximum number of General Purpose Standard instances that can be used for WorkSpaces Pools in this account in the current Region. | Yes |
| General Purpose Performance streaming instances for WorkSpaces Pools | 10 | The maximum number of General Purpose Performance streaming instances that can be used for WorkSpaces Pools in this account in the current Region. | Yes |

| Resource | Default | Description | Adjustable |
| --- | --- | --- | --- |
| General Purpose Power streaming instances for WorkSpaces Pools" | 10 | The maximum number of General Purpose Power streaming instances that can be used for WorkSpaces Pools in this account in the current Region. | Yes |
| General Purpose PowerPro streaming instances for WorkSpaces Pools" | 10 | The maximum number of General Purpose PowerPro streaming instances that can be used for WorkSpaces Pools in this account in the current Region. | Yes |
| Graphics.g4dn xlarge streaming instances for WorkSpaces Pools | 0 | The maximum number of Graphics.g4dn xlarge streaming instances that can be used for WorkSpaces Pools in this account in the current Region. | No |
| Graphics.g4dn 4xlarge streaming instances for WorkSpaces Pools | 0 | The maximum number of Graphics.g4dn 4xlarge streaming instances that can be used for WorkSpaces Pools in this account in the current Region. | No |

# Amazon WorkSpaces Instances endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

**Note**

The AWS Regions in the following table apply to WorkSpaces Instances.

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (N. Virginia) | us-east-1 | workspaces-instances.us-east-1.api.aws<br><br>workspaces-instances-fips.us-east-1.api.aws | HTTPS |
| US East (Ohio) | us-east-2 | workspaces-instances.us-east-2.api.aws<br><br>workspaces-instances-fips.us-east-2.api.aws | HTTPS |
| US West (Oregon) | us-west-2 | workspaces-instances.us-west-2.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| | | workspaces-instances-fips.us-west-2.api.aws | |
| Africa (Cape Town) | af-south-1 | workspaces-instances.af-south-1.api.aws | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | workspaces-instances.ap-east-1.api.aws | HTTPS |
| Asia Pacific (Malaysia) | ap-southeast-5 | workspaces-instances.ap-southeast-5.api.aws | HTTPS |
| Asia Pacific (Mumbai) | ap-south-1 | workspaces-instances.ap-south-1.api.aws | HTTPS |
| Asia Pacific (Seoul) | ap-northeast-2 | workspaces-instances.ap-northeast-2.api.aws | HTTPS |
| Asia Pacific (Singapore) | ap-southeast-1 | workspaces-instances.ap-southeast-1.api.aws | HTTPS |
| Asia Pacific (Sydney) | ap-southeast-2 | workspaces-instances.ap-southeast-2.api.aws | HTTPS |
| Asia Pacific (Tokyo) | ap-northeast-1 | workspaces-instances.ap-northeast-1.api.aws | HTTPS |
| Canada (Central) | ca-central-1 | workspaces-instances.ca-central-1.api.aws | HTTPS |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| Europe (Frankfurt) | eu-central-1 | workspaces-instances.eu-central-1.api.aws | HTTPS |
| Europe (Ireland) | eu-west-1 | workspaces-instances.eu-west-1.api.aws | HTTPS |
| Europe (London) | eu-west-2 | workspaces-instances.eu-west-2.api.aws | HTTPS |
| Europe (Paris) | eu-west-3 | workspaces-instances.eu-west-3.api.aws | HTTPS |
| Europe (Spain) | eu-south-2 | workspaces-instances.eu-south-2.api.aws | HTTPS |
| Middle East (UAE) | me-central-1 | workspaces-instances.me-central-1.api.aws | HTTPS |
| Israel (Tel Aviv) | il-central-1 | workspaces-instances.il-central-1.api.aws | HTTPS |
| South America (São Paulo) | sa-east-1 | workspaces-instances.sa-east-1.api.aws | HTTPS |

## Service quotas

| Resource | Default | Description | Adjustable |
|---|---|---|---|
| WorkSpaces Core Managed Instances | 1000 | The maximum number of WorkSpaces Managed Instances in this | Yes |

| Resource | Default | Description | Adjustable |
| --- | --- | --- | --- |
| | | account in the current Region. | |
| Concurrent WorkSpaces Managed Instances in allocating state | 700 | The maximum number of concurrent WorkSpaces Managed Instances in allocating state in this account in the current Region. | No |

# Amazon WorkSpaces Secure Browser endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4 endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol | |
| --- | --- | --- | --- | --- |
| US East (N. Virginia) | us-east-1 | workspaces-web.us-east-1.amazonaws.com<br>workspaces-web-fips.us-east-1.amazonaws.com | HTTPS<br>HTTPS | |
| US West (Oregon) | us-west-2 | workspaces-web.us-west-2.amazonaws.com<br>workspaces-web-fips.us-west-2.amazonaws.com | HTTPS<br>HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Mumbai) | ap-south-1 | workspaces-web.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (Singapore) | ap-southeast-1 | workspaces-web.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | workspaces-web.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | workspaces-web.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | workspaces-web.ca-central-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | workspaces-web.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | workspaces-web.eu-west-1.amazonaws.com | HTTPS | |
| Europe (London) | eu-west-2 | workspaces-web.eu-west-2.amazonaws.com | HTTPS | |

# Service quotas

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Maximum concurrent sessions for the `standard.regular` instance type in the specified region | Each supported Region: 25 | Yes | The maximum number of concurrent sessions that you can use for the `standard.regular` instance type in this account in the specified Region. |
| Maximum concurrent sessions for the `standard.large` instance type in the specified region | Each supported Region: 10 | Yes | The maximum number of concurrent sessions that you can use for the `standard.large` instance type in this account in the specified Region. |
| Maximum concurrent sessions for the `standard.xlarge` instance type in the specified region | Each supported Region: 5 | Yes | The maximum number of concurrent sessions that you can use for the `standard.xlarge` instance type in this account in the specified Region. |
| Number of web portals | Each supported Region: 3 | Yes | The maximum number of Amazon WorkSpaces Secure Browser portals in |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| | | | this account in the current Region. |
| Number of browser settings | Each supported Region: 5 | Yes | The maximum number of Amazon WorkSpaces Secure Browser browser settings in this account in the current Region. |
| Number of certificates per trust store | Each supported Region: 100 | Yes | The maximum number of certificates in an Amazon WorkSpaces Secure Browser trust store in this account in the current Region. |
| Number of identity providers per web portal | Each supported Region: 1 | Yes | The maximum number of identity providers for an Amazon WorkSpaces Secure Browser portal in this account in the current Region. |
| Number of network settings | Each supported Region: 5 | Yes | The maximum number of Amazon WorkSpaces Secure Browser network settings in this account in the current Region. |

| Name | Default | Adjustable | Description |
| --- | --- | --- | --- |
| Number of trust stores | Each supported Region: 5 | Yes | The maximum number of Amazon WorkSpaces Secure Browser trust stores in this account in the current Region. |
| Number of user settings | Each supported Region: 5 | Yes | The maximum number of Amazon WorkSpaces Secure Browser user settings in this account in the current Region. |
| Number of IP access settings | Each supported Region: 5 | Yes | The maximum number of Amazon WorkSpaces Secure Browser IP access settings in this account in the current Region. |
| Number of user access logging settings | Each supported Region: 5 | Yes | The maximum number of Amazon WorkSpaces Secure Browser user access logging settings in this account in the current Region. |

## AWS X-Ray endpoints and quotas

To connect programmatically to an AWS service, you use an endpoint. AWS services offer the following endpoint types in some or all of the AWS Regions that the service supports: IPv4

endpoints, dual-stack endpoints, and FIPS endpoints. Some services provide global endpoints. For more information, see AWS service endpoints.

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see AWS service quotas.

The following are the service endpoints and service quotas for this service.

## Service endpoints

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| US East (Ohio) | us-east-2 | xray.us-east-2.amazonaws.com | HTTPS |
| | | xray-fips.us-east-2.amazonaws.com | HTTPS |
| US East (N. Virginia) | us-east-1 | xray.us-east-1.amazonaws.com | HTTPS |
| | | xray-fips.us-east-1.amazonaws.com | HTTPS |
| US West (N. California) | us-west-1 | xray.us-west-1.amazonaws.com | HTTPS |
| | | xray-fips.us-west-1.amazonaws.com | HTTPS |
| US West (Oregon) | us-west-2 | xray.us-west-2.amazonaws.com | HTTPS |
| | | xray-fips.us-west-2.amazonaws.com | HTTPS |
| Africa (Cape Town) | af-south-1 | xray.af-south-1.amazonaws.com | HTTPS |
| Asia Pacific (Hong Kong) | ap-east-1 | xray.ap-east-1.amazonaws.com | HTTPS |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Hyderabad) | ap-south-2 | xray.ap-south-2.amazonaws.com | HTTPS | |
| Asia Pacific (Jakarta) | ap-southeast-3 | xray.ap-southeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Malaysia) | ap-southeast-5 | xray.ap-southeast-5.amazonaws.com | HTTPS | |
| Asia Pacific (Melbourne) | ap-southeast-4 | xray.ap-southeast-4.amazonaws.com | HTTPS | |
| Asia Pacific (Mumbai) | ap-south-1 | xray.ap-south-1.amazonaws.com | HTTPS | |
| Asia Pacific (New Zealand) | ap-southeast-6 | xray.ap-southeast-6.amazonaws.com | HTTPS | |
| Asia Pacific (Osaka) | ap-northeast-3 | xray.ap-northeast-3.amazonaws.com | HTTPS | |
| Asia Pacific (Seoul) | ap-northeast-2 | xray.ap-northeast-2.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Asia Pacific (Singapore) | ap-southeast-1 | xray.ap-southeast-1.amazonaws.com | HTTPS | |
| Asia Pacific (Sydney) | ap-southeast-2 | xray.ap-southeast-2.amazonaws.com | HTTPS | |
| Asia Pacific (Taipei) | ap-east-2 | xray.ap-east-2.amazonaws.com | HTTPS | |
| Asia Pacific (Thailand) | ap-southeast-7 | xray.ap-southeast-7.amazonaws.com | HTTPS | |
| Asia Pacific (Tokyo) | ap-northeast-1 | xray.ap-northeast-1.amazonaws.com | HTTPS | |
| Canada (Central) | ca-central-1 | xray.ca-central-1.amazonaws.com | HTTPS | |
| Canada West (Calgary) | ca-west-1 | xray.ca-west-1.amazonaws.com | HTTPS | |
| Europe (Frankfurt) | eu-central-1 | xray.eu-central-1.amazonaws.com | HTTPS | |
| Europe (Ireland) | eu-west-1 | xray.eu-west-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol | |
|---|---|---|---|---|
| Europe (London) | eu-west-2 | xray.eu-west-2.amazonaws.com | HTTPS | |
| Europe (Milan) | eu-south-1 | xray.eu-south-1.amazonaws.com | HTTPS | |
| Europe (Paris) | eu-west-3 | xray.eu-west-3.amazonaws.com | HTTPS | |
| Europe (Spain) | eu-south-2 | xray.eu-south-2.amazonaws.com | HTTPS | |
| Europe (Stockholm) | eu-north-1 | xray.eu-north-1.amazonaws.com | HTTPS | |
| Europe (Zurich) | eu-central-2 | xray.eu-central-2.amazonaws.com | HTTPS | |
| Israel (Tel Aviv) | il-central-1 | xray.il-central-1.amazonaws.com | HTTPS | |
| Mexico (Central) | mx-central-1 | xray.mx-central-1.amazonaws.com | HTTPS | |
| Middle East (Bahrain) | me-south-1 | xray.me-south-1.amazonaws.com | HTTPS | |
| Middle East (UAE) | me-central-1 | xray.me-central-1.amazonaws.com | HTTPS | |

| Region Name | Region | Endpoint | Protocol |
|---|---|---|---|
| South America (São Paulo) | sa-east-1 | xray.sa-east-1.amazonaws.com | HTTPS |
| AWS GovCloud (US-East) | us-gov-east-1 | xray.us-gov-east-1.amazonaws.com<br>xray-fips.us-gov-east-1.amazonaws.com | HTTPS<br>HTTPS |
| AWS GovCloud (US-West) | us-gov-west-1 | xray.us-gov-west-1.amazonaws.com<br>xray-fips.us-gov-west-1.amazonaws.com | HTTPS<br>HTTPS |

## Service quotas

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Custom sampling rules per region | Each supported Region: 25 | Yes | The maximum number of custom sampling rules per region. |
| Groups in an account | Each supported Region: 25 | No | The maximum number of groups per account. |
| Indexed annotations per trace | Each supported Region: 50 | No | The maximum number of annotations that can be indexed within a single trace. |
| Segment document size | Each supported Region: 64 Kilobytes | No | The maximum size for segment documents. |

| Name | Default | Adjustable | Description |
|---|---|---|---|
| Segments per second | Each supported Region: 2,600 | No | The maximum number of segments per second you can send to X-Ray. |
| Tags per custom sampling rule | Each supported Region: 50 | No | The maximum number of tags per custom sampling rule. |
| Tags per group | Each supported Region: 50 | No | The maximum number of tags per group. |
| Trace and service graph retention in days | Each supported Region: 30 | No | The number of days to retain trace and service map data. |
| Trace data modification period in days | Each supported Region: 7 | No | The number of days to update recorded data at no additional cost. |
| Trace document size (dynamic upper limit) | Each supported Region: 500 Kilobytes | No | The maximum size of a trace document. |
| Trace document size (lower limit) | Each supported Region: 100 Kilobytes | No | The maximum size of a trace document. |

# AWS Lifecycle Changes

When the availability of an AWS service or feature changes, we provide customers guidance on available alternatives and support for migration so that disruptions to customer operations are minimized.

We understand that changes in availability can impact your operations. For specific guidance, consult the relevant service documentation or contact AWS Support.

The following terms describe services and features at or nearing the end of their lifecycle:

- **Maintenance:** Customers can't on-board to these services and features. Customers already using these services and features can continue to do so. AWS will continue to operate and support these services and features but won't enhance or add functionality to them. We recommend that customers learn about and explore the alternatives suggested in the product pages and documentation.
- **Sunset:** Customers already using these services and features should plan to migrate to alternatives as recommended in the product pages and documentation. Sunset services have a sunset time line (typically 12 months). The sunset date is the date AWS will end operations and support of the service.
- **Full Shutdown:** Full shutdown represents the final stage in the lifecycle of a service or feature. Services and features in full shutdown are completely removed from the AWS portfolio and are no longer available or supported in any capacity. This state is reached after careful planning and execution of the preceding phases so that there is minimal disruption to customer operations.

## Services in Maintenance

Customers can't on-board to services and features in the maintenance stage. Customers already using these services and features can continue to do so. AWS will continue to operate and support these services and features but won't enhance or add functionality to them. We recommend that customers learn about and explore the alternatives suggested in the product pages and documentation.

| Service | Announcement Date | Resources |
| --- | --- | --- |
| AWS Application Discovery Service | October 7, 2025 | AWS Application Discovery Service availability change |
| AWS Cloud9 | July 25, 2024 | How to migrate from AWS Cloud9 to AWS IDE Toolkits or AWS CloudShell |
| Amazon Cloud Directory | October 7, 2025 | Amazon Cloud Directory availability change |
| Amazon CloudSearch | July 25, 2024 | Transition from Amazon CloudSearch to Amazon OpenSearch Service |
| Amazon CodeCatalyst | November 7, 2025 | How to migrate from Amazon CodeCatalyst |
| Amazon CodeGuru Reviewer | October 7, 2025 | Amazon CodeGuru Reviewer availability change |
| Amazon Comprehend - Topic Modeling, Event Detection, and Prompt Safety Classification | March 31, 2026 | Amazon Comprehend - Topic Modeling, Event Detection, and Prompt Safety Classification availability change guide |
| Amazon Forecast | July 29, 2024 | Transition your Amazon Forecast usage to Amazon SageMaker Canvas |
| Amazon FSx File Gateway | October 28, 2024 | Switch your file share access from Amazon FSx File Gateway to Amazon FSx for Windows File Server |
| AWS IoT FleetWise | March 31, 2026 | AWS IoT FleetWise availability change guide |
| AWS IoT SiteWise Data Processing Pack | October 7, 2025 | Data processing pack availability change |

| Service | Announcement Date | Resources |
| --- | --- | --- |
| AWS IoT SiteWise Monitor | October 7, 2025 | AWS IoT SiteWise Monitor availability change |
| AWS Mainframe Modernization Service (Managed Runtime Environment experience) | November 7, 2025 | AWS Mainframe Modernization availability change |
| AWS Migration Hub | November 7, 2025 | AWS Migration Hub availability change |
| Amazon Monitron | October 31, 2024 | Maintain access and consider alternatives for Amazon Monitron |
| AWS HealthOmics variant and annotation store | November 7, 2025 | AWS HealthOmics variant store and annotation store availability change |
| Amazon Rekognition - Streaming Events and Batch Image Content Moderation | March 31, 2026 | Streaming Events availability change guide and Batch Image Content Moderation availability change guide |
| Amazon Application Recovery Controller (ARC) - Readiness Check | March 31, 2026 | Amazon Application Recovery Controller (ARC) - Readiness Check availability change guide |
| Amazon S3 Glacier Direct | October 7, 2025 | Guidance for data transfer from Amazon S3 Glacier Vaults to Amazon S3 |
| Amazon S3 Glacier Select | July 25, 2024 | How to optimize querying your data in Amazon S3 |
| Amazon S3 Object Lambda | November 7, 2025 | Amazon S3 Object Lambda availability change |

| Service | Announcement Date | Resources |
| --- | --- | --- |
| Amazon S3 Select | July 25, 2024 | How to optimize querying your data in Amazon S3 |
| Amazon Simple Notification Service (Amazon SNS) - Message Data Protection (MDP) | March 31, 2026 | Amazon Simple Notification Service (Amazon SNS) - Message Data Protection (MDP) availability change guide |
| AWS Snowball Edge Compute Optimized | October 7, 2025 | AWS Snowball Edge availability change |
| AWS Snowball Edge Storage Optimized | October 7, 2025 | AWS Snowball Edge availability change |
| AWS Systems Manager Incident Manager | November 7, 2025 | AWS Systems Manager Incident Manager availability change |
| AWS Systems Manager Change Manager | November 7, 2025 | AWS Systems Manager Change Manager availability change |
| AWS App Runner | March 31, 2026 | AWS App Runner availability change guide |
| AWS Audit Manager | March 31, 2026 | AWS Audit Manager availability change guide |
| AWS CloudTrail - Lake | March 31, 2026 | AWS CloudTrail - Lake availability change guide |
| AWS Glue - Ray jobs | March 31, 2026 | AWS Glue - Ray jobs availability change guide |
| AWS .NET Modernization Tools | November 7, 2025 | AWS .NET Modernization Tools availability change |

| Service | Announcement Date | Resources |
| --- | --- | --- |
| AWS Thinkbox Deadline 10 | November 7, 2025 | AWS Thinkbox Deadline 10 |
| Amazon Timestream for LiveAnalytics | June 20, 2025 | What is Timestream for InfluxDB? |
| AWS Usage Report | October 7, 2025 | AWS Usage Report documentation |
| Amazon WorkSpaces PCoIP Web Access | November 7, 2025 | WorkSpaces Web Access |
| Amazon Bedrock Agents Classic | June 30, 2026 | Amazon Bedrock Agents Classic availability change |
| Amazon Cognito Sync | June 30, 2026 | Amazon Cognito Sync availability change |
| Amazon Kendra | June 30, 2026 | Amazon Kendra availability change |
| Amazon Q Business | June 30, 2026 | Amazon Q Business availability change |
| AWS Directory Service – Simple AD | June 30, 2026 | AWS Directory Service – Simple AD availability change |
| AWS IoT Device Defender – Detect | June 30, 2026 | AWS IoT Device Defender – Detect availability change |
| AWS Mainframe Modernization – Self-Managed Experience | June 30, 2026 | AWS Mainframe Modernization – Self-Managed Experience availability change |
| AWS Management Console – myApplications | June 30, 2026 | AWS Management Console – myApplications availability change |
| AWS Resource Groups – Group Lifecycle Events | June 30, 2026 | AWS Resource Groups – Group Lifecycle Events availability change |

| Service | Announcement Date | Resources |
| --- | --- | --- |
| AWS Service Catalog – Application Registry | June 30, 2026 | AWS Service Catalog – Application Registry availability change |
| AWS Systems Manager – Application Manager | June 30, 2026 | AWS Systems Manager – Application Manager availability change |
| Amazon SageMaker AI – A2I | June 30, 2026 | Amazon SageMaker AI – A2I availability change |
| Amazon SageMaker AI – Clarify | June 30, 2026 | Amazon SageMaker AI – Clarify availability change |
| Amazon SageMaker AI – Debugger | June 30, 2026 | Amazon SageMaker AI – Debugger availability change |
| Amazon SageMaker AI – GeoSpatial | June 30, 2026 | Amazon SageMaker AI – GeoSpatial availability change |
| Amazon SageMaker AI – Ground Truth | June 30, 2026 | Amazon SageMaker AI – Ground Truth availability change |
| Amazon SageMaker AI – Model Monitor | June 30, 2026 | Amazon SageMaker AI – Model Monitor availability change |
| Amazon SageMaker AI – Role Manager | June 30, 2026 | Amazon SageMaker AI – Role Manager availability change |
| Amazon SageMaker AI – Studio Lab | June 30, 2026 | Amazon SageMaker AI – Studio Lab availability change |

## Services in Sunset

Customers already using services and features in the sunset stage should plan to migrate to alternatives as recommended in the product pages and documentation. Sunset services have a

sunset time line (typically 12 months). The sunset date is the date AWS will end operations and support of the service.

| Service | Announcement Date | End of Support Date | Resources |
|---|---|---|---|
| Amazon FinSpace | October 7, 2025 | October 7, 2026 | Amazon FinSpace end of support |
| Amazon Fraud Detector | October 7, 2025 | October 7, 2026 | Amazon Fraud Detector availability change |
| AWS IoT Greengrass V1 | October 7, 2025 | October 7, 2026 | Migrate from AWS IoT Greengrass Version 1 |
| Amazon Lookout for Equipment | October 7, 2025 | October 7, 2026 | Preserve access and explore alternatives for Amazon Lookout for Equipment |
| Amazon Pinpoint | May 20, 2025 | October 30, 2026 | Amazon Pinpoint end of support |
| Amazon Q Developer IDE plugins | April 30, 2026 | April 30, 2027 | Amazon Q Developer IDE plugins end of support |
| Amazon RDS Custom for Oracle | March 31, 2026 | March 31, 2027 | Amazon RDS Custom for Oracle availability guide |
| AWS AppMesh | September 24, 2024 | September 30, 2026 | Migrating from AWS AppMesh to Amazon ECS Service Connect |
| AWS Proton | October 7, 2025 | October 7, 2026 | AWS Proton Service Deprecation and Migration Guide |
| AWS Service Management Connector | March 31, 2026 | March 31, 2027 | AWS Service Management Connector availability change guide |

| Service | Announcement Date | End of Support Date | Resources |
|---|---|---|---|
| AWS WAF Classic | October 7, 2025 | October 7, 2026 | AWS WAF Classic documentation |
| Amazon WorkMail | March 31, 2026 | March 31, 2027 | Amazon WorkMail availability change guide |
| Amazon WorkSpaces - Thin Client | March 31, 2026 | March 31, 2027 | Amazon WorkSpaces - Thin Client availability change guide |
| Amazon Simple Notification Service – Email Subscriptions (China Regions) | June 30, 2026 | June 30, 2027 | Amazon Simple Notification Service – Email Subscriptions (China Regions) availability change |
| Amazon WorkSpaces – PCoIP | June 30, 2026 | June 30, 2027 | Amazon WorkSpaces – PCoIP end of support |
| Amazon WorkSpaces – Pool | June 30, 2026 | June 30, 2027 | Amazon WorkSpaces – Pool end of support |
| AWS Direct Connect (Shenzhen Region) | June 30, 2026 | June 30, 2027 | AWS Direct Connect (Shenzhen Region) availability change |
| AWS Managed Services (AMS) Advanced | June 30, 2026 | June 30, 2027 | AWS Managed Services (AMS) Advanced sunset plan |
| AWS re:Post Private | June 30, 2026 | June 30, 2027 | AWS re:Post Private end of support |
| Amazon SageMaker AI – Profiler | June 30, 2026 | June 30, 2027 | Amazon SageMaker AI – Profiler availability change |

| Service | Announcement Date | End of Support Date | Resources |
|---|---|---|---|
| Amazon SageMaker AI – Mechanical Turk | June 30, 2026 | September 30, 2026 | Amazon SageMaker AI – Mechanical Turk availability change |

## Services in Full Shutdown

Full shutdown represents the final stage in the lifecycle of a service or feature. Services and features in full shutdown are completely removed from the AWS portfolio and are no longer available or supported in any capacity. This state is reached after careful planning and execution of the preceding phases so that there is minimal disruption to customer operations.

| Service | End of Support Date |
|---|---|
| AWS Application Discovery Service (Discovery Connector) | November 17, 2025 |
| AWS Backup Storage Service | June 12, 2024 |
| AWS BugBust | August 13, 2025 |
| Amazon Chime SDK - Proxy Sessions | March 31, 2026 |
| CloudEndure DR - GovCloud (CEDR) | October 1, 2025 |
| Amazon CloudWatch Evidently | October 17, 2025 |
| AWS CodeStar | July 25, 2024 |
| AWS DataSync Discovery | May 20, 2025 |
| AWS DeepComposer | September 17, 2025 |
| AWS DeepLens | January 31, 2024 |
| AWS Elastic Disaster Recovery (CloudEndure DR - GovCloud) | April 24, 2025 |

| Service | End of Support Date |
| --- | --- |
| AWS Elemental MediaStore | November 12, 2025 |
| Amazon Elastic Transcoder | November 13, 2025 |
| AWS IoT Analytics | December 15, 2025 |
| AWS IoT 1-Click | January 29, 2025 |
| Fleet Hub for AWS IoT Device Management | October 18, 2025 |
| AWS IoT RoboRunner | March 14, 2024 |
| AWS IoT Things Graph | November 9, 2022 |
| Amazon Kinesis Data Analytics for SQL | January 27, 2026 |
| Governed Tables in AWS Lake Formation | December 31, 2024 |
| Amazon Lookout for Metrics | October 10, 2025 |
| Amazon Lookout for Vision | October 31, 2025 |
| Amazon Lumberyard | May 1, 2021 |
| AWS Mainframe Modernization Application Testing | October 7, 2025 |
| AWS Mobile Hub | October 31, 2021 |
| AWS NICE EnginFrame | September 25, 2025 |
| Amazon Nimble Studio | June 30, 2024 |
| AWS OpsWorks | May 1, 2024 |
| AWS OpsWorks for Chef Automate | May 1, 2024 |
| AWS OpsWorks for Puppet Enterprise | May 1, 2024 |
| AWS OpsWorks Stacks | May 1, 2024 |

| Service | End of Support Date |
| --- | --- |
| AWS Private 5G | May 20, 2025 |
| Amazon Quantum Ledger Database (Amazon QLDB) | July 31, 2025 |
| Amazon RDS on VMWare | May 1, 2022 |
| Amazon Rekognition Person Pathing | October 31, 2025 |
| AWS RoboMaker | September 10, 2025 |
| AWS Server Migration Service | April 1, 2023 |
| AWS SimSpace Weaver | March 31, 2026 |
| AWS Snowball Edge Storage Optimized 80TB (V3S) | November 12, 2025 |
| AWS Snowball Edge with GPU (V3G) | November 12, 2025 |
| AWS Snowblade | November 12, 2024 |
| AWS Snowball Edge | November 12, 2025 |
| AWS Snowmobile | March 14, 2024 |
| Amazon Sumerian | March 1, 2023 |
| Amazon WorkDocs | April 25, 2025 |
| Amazon WorkLink | November 30, 2021 |
| Amazon WorkSpaces Application Manager | September 1, 2023 |
| Amazon Connect Customer Voice ID | May 20, 2026 |
| AWS Database Migration Service Fleet Advisor | May 20, 2026 |
| AWS IoT Events | May 20, 2026 |
| AWS IQ | May 20, 2026 |

| Service | End of Support Date |
| --- | --- |
| AWS Panorama | May 20, 2026 |
| Amazon Inspector Classic | May 20, 2026 |
| AWS Elemental MediaLive and MediaPackage – Isolated Regions | June 30, 2026 |
| Amazon Chime SDK – Carrier Voice Focus | June 30, 2026 |
| Amazon SageMaker AI – Ground Truth Plus | June 30, 2026 |

# AWS Lifecycle documentation history

The following table describes updates to the AWS Lifecycle documentation. To receive notifications for changes to the AWS Lifecycle documentation, subscribe to the RSS feed.

| Change | Description | Date |
| --- | --- | --- |
| AWS Lifecycle Changes | Information about lifecycle changes for AWS services is available. Review this information to see if an AWS service is in maintenance, sunset, or full shutdown. | March 23, 2026 |