FBI FLASH
ACTIONABLE CYBER INTELLIGENCE
13 NOVEMBER 2026
FLASH Number
FLASH-20261113-001

Threat Overview of ARTEX Autonomous Penetration Testing Framework

Summary

The Federal Bureau of Investigation (FBI) is releasing this FLASH to highlight the tactics, techniques, and procedures (TTPs) and indicators of compromise (IOCs) associated with the ARTEX autonomous penetration testing framework. ARTEX is an open-source, large language model (LLM)-driven agentic platform that threat actors have weaponized to automate multiple phases of cyber intrusions at machine speed, from reconnaissance through data exfiltration. Since at least August 2026, a variety of malicious cyber actors have used ARTEX to conduct campaigns targeting organizations globally across multiple sectors, including financial services, government, technology, defense, semiconductor, energy, healthcare, manufacturing, telecommunications, and legal industries.[1,2,3,4,5]

ARTEX significantly reduces the time required to chain attack techniques, with some intrusions progressing from initial targeting to privileged access in under three hours.[2] The framework enables threat actors with minimal technical expertise to conduct complex, multi-target operations simultaneously, fundamentally accelerating attack velocity and expanding operational scale. Notable campaigns include financially motivated attacks against South Korean financial institutions resulting in exfiltration of personal information from approximately 67,000 individuals,[6,7,8] and a supply chain compromise affecting ConnectWise Cloud Backup customers that enabled access to more than 2,250 downstream tenants.[2]

Following public disclosure of malicious usage, the original ARTEX developer removed the GitHub repository in October 2026.[3] However, the framework continues to proliferate through community forks, localized variations, archived copies, and a commercial version (ARTEX PRO), ensuring continued availability to threat actors.[3,10] As of early November 2026, researchers identified over 1,600 publicly exposed ARTEX console instances, with total deployments likely significantly higher when including systems behind firewalls or in private networks.[4]

The FBI encourages organizations—particularly those operating internet-exposed applications, self-service portals, cloud infrastructure, and partner/supplier systems—to implement the recommendations in the Mitigations section to reduce the likelihood and impact of ARTEX-enabled intrusions.



Threat

ARTEX is an autonomous, multi-agent penetration testing platform written in Go.[3,4] The framework translates high-level tactical goals described in natural language into actionable machine commands executed by specialized super intelligence (SI) agents, requiring minimal human interaction beyond initial objective setting.[2,3] ARTEX does not include its own SI model but integrates with external LLM providers, including frontier models (Anthropic Claude, OpenAI GPT) and open-source or Chinese-origin models (DeepSeek, GLM, Grok, Kimi).[3,6,7,9]

Analysis of attacker-controlled ARTEX instances shows Chinese models are most prevalent, with DeepSeek Flash family models dominating usage due to their balance of capabilities, relatively low cost, and lack of observable guardrails when framed for penetration testing.[7,9]

Originally released on GitHub by developer "Autumn-27" in July 2026, ARTEX gained rapid adoption, reaching an estimated 1,500–2,000 GitHub stars within 10 weeks.[4] The developer stated the tool was intended for authorized security testing within defined scopes and made the repository unavailable on October 8, 2026, following reports of threat actor abuse.[3,6] Despite this action, ARTEX remains widely accessible through independent routes and mirrors on code-sharing platforms.[3,4,9]

Observed Capabilities and Potential Uses

Based on observed capabilities and ongoing threat actor adoption, ARTEX presents the following risks:

- **Accelerated Multi-Target Operations:** SI agents provide automation for the time-consuming phases of enumeration and exploitation; small threat groups can run concurrent intrusion campaigns at a scale previously unavailable to them.
- **Exploitation of Overlooked Attack Surfaces**: ARTEX excels at identifying and exploiting peripheral systems receiving less security attention than flagship applications, including partner and supplier portals, mobile backends, employee applications, sales-support systems, legacy services, third-party integrations, and applications created across different business units.[1]
- **Supply Chain Compromise**: Agents have demonstrated capability to identify and exploit supply chain attack vectors, including modifying iOS application distribution infrastructure via exposed cloud credentials.[9] Autonomous agents can systematically map trust relationships between organizations. By identifying and compromising weaker third-party vendors, CI/CD pipelines, or integrated software providers, attackers can leverage ARTEX to pivot from a peripheral supplier directly into high-value upstream enterprise networks.
- **Reduced Barrier to Entry**: ARTEX lowers technical barriers for conducting complex intrusions, enabling threat actors with minimal expertise to execute sophisticated campaigns.[3,9]

- **Compressed Detection Window**: ARTEX generates high volumes of forensically noisy traffic that may trigger alerts, but attack speed provides limited time for coordinated defender response before objective completion.[2,3]

Observed Tactics, Techniques, and Procedures

The following table maps observed ARTEX activity to MITRE ATT&CK phases:

Phase	Observed Activity	ATT&CK Techniques
Initial Recon	• Automatically mapped effective permissions of Azure credentials (SAS tokens, Service Principals, Kudu credentials, Storage Account keys) using sub-second capability checks, test blob writes/deletes, ARM deployment validation dry-runs • Tested compromised IAM user permissions using DryRun flags • Fuzzing web directories across 500+ requests • WHOIS asset mapping and active vulnerability scanning • Queried api.db-ip[.]com to determine victim geolocation	T1595 – Active Scanning T1580 – Cloud Infrastructure Discovery T1087.004 – Account Discovery: Cloud Account T1069.003 – Permission Groups Discovery: Cloud Groups T1526 – Cloud Service Discovery
Initial Access	• Exploited unsigned SAML assertions • Exploited XXE flaws • Exploited anonymous file upload • Abused cross-account cloud trust relationships from an external partner tenant • Compromised third-party vendor credentials to pivot via trusted B2B federation • Exploited SSRF • Hardcoded administrator	T1190 – Exploit Public-Facing Application T1078 – Valid Accounts T1133 – External Remote Services T1195 – Supply Chain Compromise T1199 – Trusted Relationship T1606.002 – Forge Web Credentials: SAML Tokens T1552.001 – Unsecured Credentials: Credentials in Files

	passwords in SPA JavaScript bundles • Missing authentication on Web APIs • RCE through command injection • Insecure Direct Object References (IDOR) • Weak administrative credentials • RCE through SSTI • SQL injection • Authentication bypass • Used stolen VPN configurations and leaked credentials	
Establish Foothold / Persistence	• Deployed JSP and PHP web shells (including via MySQL log poisoning) • Abused Kudu administrative APIs to execute arbitrary commands • Established C2 on internal Windows database servers • Installed open-source RATs (WISP, vshell, AdaptixC2) • Created SSH key pairs, allocated Elastic IPs • Account takeover (password resets) • Added credentials to existing OAuth applications	T1505.003 – Server Software Component: Web Shell T1078.004 – Valid Accounts: Cloud Accounts T1098.001 – Account Manipulation: Additional Cloud Credentials T1574.001 – Hijack Execution Flow: DLL T1036.005 – Masquerading T1105 – Ingress Tool Transfer T1059 – Command and Scripting Interpreter
Privilege Escalation	• Launched cloud compute instances with high-value IAM roles • Kerberoasting to harvest service tickets • SQL injection to dump database credentials • DCSync to steal domain krbtgt key • Created accounts with	T1134.001 – Access Token Manipulation T1098.003 – Account Manipulation: Additional Cloud Roles T1552.005 – Unsecured Credentials: Cloud Instance Metadata API T1606.002 – Forge Web Credentials: SAML Tokens

	is_admin=true flag • Abused resource-based constrained delegation • Queried cloud metadata service via XXE • Forged administrator token using unrotated authentication keys	T1606 – Forge Web Credentials T1190 – Exploit Public-Facing Application T11558.003 – Steal or Forge Kerberos Tickets: Kerberoasting
Internal Recon	• Ran iterative PowerShell enumeration scripts against Domain Controllers • Enumerated over 110,000 Active Directory objects • Read LAPS passwords • Inventoried endpoint sensor coverage to identify Windows machines without CrowdStrike Falcon	T1518.001 – Software Discovery: Security Software Discovery T1555 – Credentials from Password Stores T1087.002 – Account Discovery: Domain Account T1552.006 – Group Policy Preferences T1059.001 – Command and Scripting Interpreter: PowerShell
Lateral Movement	• Attempted lateral movement to domain controllers using Pass-the-Ticket, WinRM, Kerberos relay, Pass-the-Hash (blocked by endpoint protection)	T1550.003 – Use Alternate Authentication Material: Pass the Ticket T1550.002 – Use Alternate Authentication Material: Pass the Hash T1021.006 – Remote Services: Windows Remote Management (WinRM) T1557.001 – Adversary-in-the-Middle: LLMNR/NBT-NS Poisoning and SMB Relay
Collection / Exfiltration Prep	• Interrogated backend SQL databases • Mass blob reads • Deployed custom boot scripts via EC2 instances to harvest data • Sent command output to rotating webhook.site endpoints • Read mailboxes on	T1114.002 – Email Collection: Remote Email Collection T1555.006 – Cloud Secrets Management Stores T1530 – Data from Cloud Storage T1005 – Data from Local System T1074.001 – Data Staged: Local Data Staging

	compromised mail server • Created new 100GB volumes for exfiltration preparation	T1074.002 – Data Staged: Remote Data Staging T1059 – Command and Scripting Interpreter
Command and Control	• Used residential proxy services • Routed traffic through Tor, SOCKS proxies, compromised relays • Tunneled traffic via Chisel and SoftEther • BLOODALCHEMY communicated with C2 domains	T1090 – Proxy T1090.003 – Proxy: Multi-hop Proxy T1572 – Protocol Tunneling T1071.001 – Application Layer Protocol: Web Protocols

Indicators

The FBI recommends organizations investigate or vet IOCs prior to taking action, such as blocking. The FBI recognizes that infrastructure used for malicious activity may be hosted on cloud or virtualized platforms where IP addresses can be dynamically or temporarily assigned. As a result, an address associated with malicious activity at one point in time may later be reassigned to benign services. Accordingly, these indicators should be interpreted as historically observed infrastructure within the activity window and should be corroborated with current network telemetry or additional intelligence sources when used for defensive or investigative purposes.

IP Addresses

Indicator	Context	Source
38.244.50[.]120	ARTEX host and open directory (South Korea financial campaign)	CrowdStrike[6], Censys[4]
144.172.64[.]60 **144.172.64[.]99** **144.172.64[.]128**	ConnectWise campaign; sign-ins and Graph API activity in customer environments	Microsoft[2]
187.15.88[.]186 **187.15.90[.]217** **187.15.89[.]186**	Graph API access via compromised apps (transient VPN IPs relevant October 15–16, 2026)	Microsoft[2]

101.53.80[.]20 **205.214.59[.]31** **124.155.252[.]63** **154.201.79[.]246** **23.248.249[.]90** **23.158.220[.]98** **103.248.148[.]84** **203.160.133[.]172** **209.209.85[.]38**	Proxy IP addresses (South Korea financial campaign)	CrowdStrike[6]
199.19.95[.]66 **199.19.95[.]67** **199.19.95[.]68**	WISP C2 nodes (Storm-3358 campaign; active as of November 6, 2026)	Microsoft[10]
91.92.240[.]238 **85.11.168[.]45** **178.16.54[.]129** **31.77.229[.]35**	ARTEX servers (Storm-3358 campaign; active as of November 6, 2026)	Microsoft[10]
85.11.167[.]52	C2 server (Storm-3358 campaign; active as of November 6, 2026)	Microsoft[10]

Domains

Indicator	Context	Source
xcai[.]pro	LLM API reseller (South Korea financial campaign)	CrowdStrike[6]

Network and Behavioral Indicators

Indicator Type	Description	Source
Outbound API calls	api.anthropic.com, api.deepseek.com, open.bigmodel.cn, api.x.ai (investigate unexpected calls alongside other suspicious activity)	Zscaler[7]
HTTP request patterns	High-frequency automated requests with sequential enumeration against login/search endpoints	Zscaler[7]
Data transfers	Outbound transfers of structured data (names, phone numbers, financial information) over HTTPS to	Zscaler[7]

	unauthorized external destinations	
Web shells	*.aspx or *.jsp files in unexpected web application paths	Zscaler[7]

Information Requested

The FBI encourages recipients of this document to report information concerning suspicious or criminal activity to their local FBI field office at www.fbi.gov/contact-us/field-offices or the FBI Internet Crime Complaint Center (IC3) at www.ic3.gov.

If reporting an incident involving ARTEX or similar agentic SI frameworks, please provide the following information:

1. **ARTEX Infrastructure Artifacts**: Source code or screenshots of suspected ARTEX console interfaces and ARTEX configuration files

2. **Authentication and Access Logs**: Service principal sign-in logs showing unusual source networks, ASNs, hosting providers, or geographies; user agent strings from suspicious authentication attempts; IP login information and associated timestamps; evidence of authentication from previously unseen infrastructure

3. **Application and API Activity**: Microsoft Graph activity; Microsoft Entra audit records; Microsoft 365 audit records showing reconnaissance; sequential lookups or unusually high numbers of records accessed per session on partner/employee portals; rate limit violations or volumetric anomalies in API calls

4. **LLM and SI Infrastructure Indicators**: Outbound connections to LLM API endpoints; evidence of self-hosted LLM gateway or reseller software (CLIProxyAPI, New API); unfamiliar LLM proxies or resellers in egress logs; token consumption patterns or API billing anomalies

5. **Cloud and Container Activity**: AWS IAM activity logs; Azure ARM deployment validation dry-runs across multiple resource types; test blob write/delete operations in cloud storage with sub-second timing; evidence of compromised Kudu/SCM credentials or unusual Kudu API activity; container orchestration logs; cloud metadata service queries; unexpected IAM role assumptions, PassRole operations, or instance profile modifications

6. **Exploitation and Persistence Artifacts**: Web shells; evidence of BLOODALCHEMY components; open-source RAT indicators (WISP, vshell, AdaptixC2); newly created accounts with naming patterns;

application credential modifications to OAuth applications; SSH key additions, Elastic IP allocations, unexpected EC2/Azure VM launches

7. **Attack Characteristics**: Compressed bursts of activity; high volumes of failed attempts with iterative, incremental changes; evidence of multiple sub-agents overwriting each other or generating duplicate requests; mass notification/email broadcasts from compromised applications

8. **Loss and Impact Information**: Total number of compromised accounts or affected users; types and volumes of data accessed or exfiltrated; time from initial suspicious activity to detection; time from detection to containment; total attempted loss amounts

Recommendations

The FBI recommends organizations take the steps below to improve their security posture in response to ARTEX and similar agentic SI frameworks.

Immediate Actions

Reduce External Attack Surface:

• Inventory all internet-facing applications, including peripheral systems: partner and supplier portals, mobile backends, employee applications, sales-support systems, legacy services, third-party integrations, and applications created across different business units.[1,11]

• Restrict cloud administration, network device management, application publishing interfaces (Kudu/SCM endpoints), hypervisors, backup platforms, and other sensitive services to trusted networks, private connectivity, or dedicated administrative paths.[5]

• Disable basic publishing authentication where not required; restrict access to deployment and publishing credentials; restrict SCM access to approved management networks or deployment infrastructure.[5]

• Consider deployment models such as run-from-package that reduce direct modification of deployed application files.[5]

• Establish governance around public IP allocation; require architecture review before workloads or management interfaces are exposed directly to the internet.[5]

• Move enterprise applications behind Zero Trust access solutions to eliminate external exposure where feasible.[7]

Harden Identity and Access:

• Apply least privilege to multi-tenant application permissions and consent grants.[2]

• Monitor credential additions, privilege assignments, and anomalous app-only activity, especially to applications with high privileged permissions and scopes.[2]

• Review every Microsoft 365 tenant where affected ConnectWise applications are currently or were previously installed or authorized, even if integration no longer in use; search by application ID rather than display name.[2]

• Review successful and failed service-principal sign-ins separately; investigate unfamiliar sources and activity outside expected vendor infrastructure (do not limit review to published IPs).[2]

• Review Microsoft Entra audit records for application/service principal credential additions; new applications/service principals/owners; app-role assignments and consent grants; privileged role assignments; changes to users/groups/authentication methods/Conditional Access (CA). If suspicious activity is confirmed, revoke unnecessary permissions and rotate credentials.[2]

• Require certificate manager or CA approval for higher-risk enrollment scenarios; enable Extended Protection for Authentication on CA web interfaces where supported; monitor unusual certificate enrollment involving privileged users or unexpected requesting systems.[5]

Strengthen Application Protections:

• Put guardrails around high-impact business APIs: rate limits, recipient limits, transaction thresholds, abuse controls appropriate to function; alert when application operates materially outside normal volume.[5]

• On partner and employee portals, alert on sequential lookups and unusually high numbers of records accessed per session; apply rate limits to affected endpoints; enforce server-side authorization for access to customer records and limit responses to fields required for the operation.[7]

• Avoid returning detailed stack traces, compiler messages, SQL errors, internal paths, framework details to unauthenticated users; keep detailed troubleshooting information in internal logs.[5]

• Review application configurations, environment variables, infrastructure-as-code, deployment scripts, CI/CD pipelines, databases for embedded secrets; implement continuous scanning of repositories, container images, build artifacts, pipelines for exposed secrets; rotate exposed credentials promptly.[5]

Preserve Evidence and Respond:

• Preserve available Microsoft Entra sign-in and audit records, Microsoft Purview Audit records, SIEM data, previously configured Microsoft Graph telemetry; do not delay urgent containment while waiting for exports.[2]

• Review activity starting from at least October 1, 2026, for ConnectWise-related activity; extend service-principal sign-in review and investigate earlier related activity including failed authentications and volume anomalies.[2]



• If suspicious activity is confirmed, have appropriately authorized responder remove confirmed unauthorized credentials, grants, owners, assignments, identities after preserving evidence.[2]

• Keep affected ConnectWise application connections disabled until ConnectWise provides formal restoration instructions; do not delete enterprise applications as first response (preserve configuration and evidence before irreversible changes).[2]



Notes

_____

1. Akamai, "What ARTEX Reveals About Overlooked Attack Surfaces," October 23, 2026.

2. Microsoft Security Research and Microsoft Threat Intelligence, "AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise," October 21, 2026.

3. Google Threat Intelligence Group, "ARTEX and Friends: Enter Agentic Offensive Frameworks," October 30, 2026.

4. Censys, "ARTEX: Agentic Asbestos," November 5, 2026.

5. Google Threat Intelligence Group and Mandiant, "ARTEX and Friends: Enter Agentic Offensive Frameworks – Defensive Strategies from Mandiant Consulting," October 30, 2026.

6. CrowdStrike, "Unknown Threat Actor Uses AI-Driven ARTEX to Target South Korean Finance," October 7, 2026.

7. Zscaler ThreatLabz, "Threat Actor Abuses Agentic AI Pentesting Framework to Target South Korean Financial Institutions," October 14, 2026.

8. CERT-EU, "From Penetration Testing to Cyber Threat: Beware of ARTEX," November 5, 2026.

9. Check Point Research, "Hacking on Autopilot: Inside ARTEX Cybercrime Campaigns," November 5, 2026.

10. Microsoft Threat Intelligence, "ARTEX abuse expands as cybercriminals target financial sector," November 2026.

11. Akamai, "What ARTEX Reveals About Overlooked Attack Surfaces – API Security Implications," October 23, 2026.

Reporting Notice

If you identify suspicious activity within your enterprise or have information related to the contents of this document, please contact your local FBI Cyber Squad immediately at **www.fbi.gov/contact-us/field-offices**. The FBI also encourages you to report suspicious or criminal activity to the FBI Internet Crime Complaint Center at **www.ic3.gov**. When available, each report should include the date, time, location, type of activity, number of people, and type of equipment used for the activity, the name of the submitting company or organization, and a designated point of contact. Press inquiries should be directed to the FBI's National Press Office at npo@fbi.gov or (202) 324-3691.

Individual indicators included in this document should always be evaluated in light of your complete information security situation. Some indicators, particularly those of a nondeterministic or ephemeral nature (such as filenames or IP addresses), may not be indicative of a compromise.

Your organization has no obligation to provide information in response to this product. If, after reviewing the information provided, your organization decides to provide information to the FBI, it must do so consistent with applicable state and federal law.

Administrative Note

The information in this document is being provided by the FBI, with no guarantees or warranties, for potential use at the sole discretion of recipients to protect against cyber threats. This data is provided to help cybersecurity professionals and system administrators guard against the persistent malicious actions of cyber actors. The FBI does not endorse any commercial entity, product, company, or service, including any entities, products, or services linked within this document. Any reference to specific commercial entities, products, processes, or services by service mark, trademark, manufacturer, or otherwise, does not constitute or imply endorsement, recommendation, or favoring by the FBI.

This FLASH is marked **TLP:CLEAR**. The information in this product may be shared without restriction. Information is subject to standard copyright rules.

