CYBER THREAT INTELLIGENCE

CRITICALSTART® Security Advisory

TLP CLEAR // [CS-SA-26-1103] ARTEX: Exploited Open-Source Penetration Platform

Executive Summary

The CRITICALSTART® Cyber Research Unit (CRU) is tracking the malicious use of ARTEX, an open-source, multi-agent penetration testing framework that uses large language models (LLMs) to automate intrusions from reconnaissance through data exfiltration.[1] Since at least August 2026, state-aligned espionage groups and financially motivated criminals have used ARTEX against organizations in financial services, government, technology, defense, semiconductor, energy, healthcare, manufacturing, telecommunications, and legal sectors.[1, 5] The framework uses no new exploits; it chains familiar weaknesses such as exposed credentials, over-permissioned cloud applications, and common web flaws at a pace that reached privileged cloud access in under three hours in some intrusions.[1, 2] On November 13, 2026, the Federal Bureau of Investigation (FBI) issued FLASH-20261113-001 warning that the developer's October removal of the source code has not stopped its spread.[1] Organizations should prioritize third-party application identity review, secret rotation, and behavior-based detection over indicator blocking.

Overview

ARTEX is a Go-based framework with an embedded Next.js web console.[5] Planning agents convert an operator's natural-language objective into tasks, worker agents run standard offensive tooling, and results feed back until the objective is met.[8, 9] The framework has no built-in model and connects to commercial or open-weight LLMs. Analysis of attacker-controlled instances shows DeepSeek Flash family models dominate because of low cost and few effective guardrails, with GLM, Grok, and Kimi also observed.[1, 6] The developer closed the source on October 8, 2026, citing misuse.[1] Forks, localized variants, archived copies, and a paid ARTEX PRO edition remain in circulation, and researchers have found backdoored ARTEX Docker images on public registries.[1, 9, 10]

Scope and Scale

The ARTEX threat landscape spans at least five publicly reported operator clusters, including China-linked espionage groups and Asia-based cybercriminal operators.[10] Its exposed infrastructure is substantial, with 1,465 internet-facing ARTEX consoles identified as of November 4, 2026. A separate assessment identified 1,608 active panels among 2,557 historically observed hosts.[6, 8] The largest reported espionage campaign targeted more than 200 organizations across at least 14 countries within 11 days, with post-compromise activity confirmed at nine organizations.[3] Separately, the largest supply chain campaign generated sign-in attempts against more than 2,250 Microsoft 365 tenants, peaking at 2,900 attempts per minute.[2] Attackers demonstrated the ability to obtain privileged cloud access in under three hours, while one suspected incident involved access to Azure Key Vault data just five minutes after a secret was exposed.[1, 10] These operations also required relatively modest resources, with reported LLM usage costs ranging from $145 to $1,500 per campaign.[6]

© 2026 Critical Start. All rights reserved.

Attribution

No public evidence links these clusters beyond shared use of the tool by the following:[11]

- VAPOR PANDA (CrowdStrike): China-nexus espionage group targeting foreign ministries, telecommunications, defense, semiconductor, and pharmaceutical organizations; separate from Storm-0940.[3]
- Storm-0940 (Microsoft; overlaps GTIG's UNC7228): China-state aligned actor behind the ConnectWise Cloud Backup supply chain compromise.[2, 5]
- Storm-3358 (Microsoft): Cambodia-based, financially motivated cluster targeting Southeast Asian financial organizations.[7]
- Unattributed South Korea actor: financially motivated Chinese speaker (moderate confidence) whose campaign exposed personal data of about 67,000 people; related activity tracked as TAG-217.[1, 4, 11]

Timeline of Key Events

Date (2026)	Event
July	ARTEX released as open source on GitHub[1]
Late August	Earliest in-the-wild use identified in incident response engagements[5]
September	Storm-0940 begins using ARTEX to automate attacks[2]
October 8	Developer halts updates and closes the source[1]
October 17	VAPOR PANDA campaign begins; ConnectWise disables seven Cloud Backup application identities[3, 11]
October 21	Microsoft attributes the ConnectWise Cloud Backup supply chain compromise to Storm-0940[2]
October 29 to November 6	CrowdStrike, GTIG/Mandiant, Censys, Check Point, and Microsoft publish campaign analyses[3, 5, 6, 7, 8]
November 13	FBI issues FLASH-20261113-001[1]

Attack Methodology

ARTEX campaigns follow a conventional intrusion lifecycle. What changes is speed, parallelism, and persistence: agents retry failed steps within seconds and pursue multiple paths at once. Google Threat Intelligence Group (GTIG) recorded 27 cloud permission checks in 23 seconds during one intrusion.[5, 10]

Initial Access

Agents enter through whatever an internet-facing system exposes. The FBI lists unsigned Security Assertion Markup Language (SAML) assertions, XML External Entity (XXE) injection, Server-Side Request Forgery (SSRF), SQL injection, command injection, Server-Side Template Injection (SSTI), Insecure Direct Object Reference (IDOR), missing application programming interface (API) authentication, anonymous file upload, and hardcoded administrator passwords in JavaScript bundles.[1] ARTEX performs best against peripheral systems that receive less security attention, including partner portals, mobile backends, sales-support applications, and legacy services.[1, 12]

Identity and Cloud Abuse

Storm-0940 obtained credentials for ConnectWise Cloud Backup application identities, which already held consented permissions in customer Microsoft 365 tenants.

© 2026 Critical Start. All rights reserved.

The applications authenticated as themselves, so no user sign-in or Multi-Factor Authentication (MFA) prompt was involved.[2, 11] Where sign-in succeeded, agents ran Graph reconnaissance, added application credentials, and assigned privileged roles.[2] Across Azure and AWS, agents also exfiltrated Key Vault data, ran scripts through VM Run Command, and searched DevOps and CloudWatch logs for secrets.[10]

On-Premises Post-Exploitation

GTIG and Mandiant observed agents enumerate more than 110,000 Active Directory objects, Kerberoast service accounts, read Local Administrator Password Solution (LAPS) passwords, and use DCSync to obtain the krbtgt key.[1, 5] Agents deployed JSP web shells, sent output to 491 rotating webhook.site URLs, and fetched the WISP, vshell, and AdaptixC2 remote access tools.[7, 10, 11] VAPOR PANDA directed its agents to find Windows hosts with no Endpoint Detection and Response (EDR) sensor, then deployed the BLOODALCHEMY backdoor through Dynamic Link Library (DLL) sideloading by a renamed GameHook binary running as svchost.exe.[3, 10]

MITRE ATT&CK Mapping

The mapping below consolidates techniques reported by the FBI, CrowdStrike, Microsoft, and GTIG/Mandiant.[1, 3, 4, 5]

ID	Name	Observed Activity
T1588.007	Obtain Capabilities: Artificial Intelligence	Acquisition and deployment of ARTEX and commercial or open-weight LLMs[4]
T1595	Active Scanning	Directory fuzzing of 500+ requests and high-frequency endpoint enumeration[1]
T1580	Cloud Infrastructure Discovery	Sub-second permission checks against Azure and AWS credentials[1]
T1190	Exploit Public-Facing Application	SQL injection, XXE, SSRF, SSTI, command injection, and IDOR against web applications[1]
T1195	Supply Chain Compromise	Compromise of a multi-tenant vendor application[1, 2]
T1199	Trusted Relationship	Use of pre-consented vendor application access to downstream tenants[1, 2]
T1078.004	Valid Accounts: Cloud Accounts	Authentication as compromised application identities[1, 11]
T1110.003	Brute Force: Password Spraying	Cross-tenant spraying peaking at 2,900 failures per minute[2]
T1552.001	Unsecured Credentials: Credentials in Files	Harvesting secrets from public repositories and web roots[1, 10]
T1552.005	Unsecured Credentials: Cloud Instance Metadata API	SAML and XXE chain used to query the metadata service[1, 11]
T1606.002	Forge Web Credentials: SAML Tokens	Abuse of unsigned SAML assertions; forged admin tokens using unrotated keys[1, 11]
T1505.003	Server Software Component: Web Shell	Numbered JSP and PHP web shells, including via MySQL log poisoning[1, 11]
T1098.001	Account Manipulation: Additional Cloud Credentials	Credentials added to existing OAuth applications[1, 2]
T1098.003	Account Manipulation: Additional Cloud Roles	Privileged role assignments after authentication[1, 2]
T1558.003	Steal or Forge Kerberos Tickets: Kerberoasting	Service ticket requests against Active Directory[1, 5]
T1003.006	OS Credential Dumping: DCSync	Replication of the krbtgt key and domain controller hashes[5, 11]

© 2026 Critical Start. All rights reserved.

ID	Name	Observed Activity
T1518.001	Software Discovery: Security Software Discovery	Inventory of EDR coverage to find unprotected Windows hosts[1, 3]
T1574.001	Hijack Execution Flow: DLL	BLOODALCHEMY sideloaded by a renamed GameHook binary[1, 3]
T1036.005	Masquerading: Match Legitimate Name or Location	GameHook binary renamed svchost.exe[1, 10]
T1105	Ingress Tool Transfer	Retrieval of WISP, vshell, and AdaptixC2 with curl and netcat[1, 7]
T1555.006	Credentials from Password Stores: Cloud Secrets Management Stores	Azure Key Vault exfiltration[1, 10]
T1530	Data from Cloud Storage	Mass blob reads and SharePoint and OneDrive downloads[1, 10]
T1114.002	Email Collection: Remote Email Collection	Mailbox access through Microsoft Graph[1, 10]
T1567.004	Exfiltration Over Web Service: Exfiltration Over Webhook	Command output sent to rotating webhook.site URLs[11]
T1090.003	Proxy: Multi-hop Proxy	Residential proxies, Tor, and SOCKS relays to rotate source addresses[1, 4]
T1572	Protocol Tunneling	Tunneling through Chisel and SoftEther[1]

Relevant Vulnerability Context

ARTEX has no Common Vulnerabilities and Exposures (CVE) identifier or patch, and the FBI cites vulnerability classes rather than specific CVEs.[1, 11] Operator infrastructure did hold exploit builds for two Linux flaws, shown with Common Vulnerability Scoring System (CVSS) scores.[10]

Implications for Organizations

Organizations must account for attack timelines measured in minutes, not days. Microsoft documented ARTEX-enabled intrusions that progressed from initial targeting to privileged cloud access in under three hours, while another investigation observed an AI agent move from exposed GitHub secrets to Azure Key Vault data in just five minutes.[1, 10] Automated agents can scan, exploit, validate credentials, and pursue multiple attack paths with limited human intervention. Alert generation alone is insufficient if security teams cannot rapidly investigate and revoke compromised access. Third-party application integrations and OAuth permissions also represent significant exposure. The ConnectWise compromise demonstrated how stolen application credentials could provide a path into more than 2,250 downstream Microsoft 365 tenants without requiring interactive user authentication.[2] Organizations should therefore treat application identities, service principals, OAuth grants, and third-party integrations as privileged access, with particular attention to excessive permissions and credentials that remain valid after a vendor incident.

ARTEX-related activity does not, by itself, establish malicious intent or identify the operator. Researchers have observed the same open-source frameworks across espionage, financially motivated operations, and legitimate penetration testing engagements.[6, 10] The growing availability of exposed consoles, modified forks, and alternative frameworks such as PentAGI and CAIRN further complicates attribution and detection. Organizations should focus on observable behaviors, including automated reconnaissance, password spraying, unauthorized OAuth application changes, credential harvesting, privilege escalation, and unusual Microsoft Graph activity, rather than relying solely on ARTEX-specific signatures. Given the low reported operating costs and ability to reuse or modify open-source tooling, CRU assesses with high confidence that adoption of agentic offensive frameworks will continue expanding through 2027.

© 2026 Critical Start. All rights reserved.

Prioritized Mitigation Strategies

To reduce the risk of ARTEX-enabled intrusions and limit the impact of automated attacks against cloud and enterprise environments, CRU recommends the following actions:

Immediate (Within 24 Hours)

- **Identify compromised third-party applications:** Search every Microsoft 365 tenant for the ConnectWise application IDs listed in Appendix A. Search by application ID rather than display name, keep affected connections disabled pending vendor guidance, and review associated service principal sign-in activity from October 1, 2026, onward.[1, 11]
- **Audit cloud identity changes:** Review Microsoft Entra ID audit logs dating back to August 1, 2026, for unauthorized application credentials, ownership changes, consent grants, and privileged role assignments. Investigate authentication errors AADSTS7000215 and AADSTS7000222, particularly when repeated failures are followed by a successful sign-in.[1, 10]
- **Hunt for ARTEX artifacts:** Search endpoint, network, and cloud telemetry for the indicators in Appendix B, including ARTEX-related user-agent strings, paths beginning with `artx`, and files matching `artex-probe-*`.[9, 10, 11]
- **Preserve forensic evidence:** Retain Microsoft Entra ID, Microsoft Purview Audit, and SIEM logs before disabling credentials or removing suspicious access. Avoid immediately deleting affected enterprise applications, as doing so may destroy evidence needed to investigate the scope of compromise.[1]

Short-Term (Within Two Weeks)

- **Eliminate exposed credentials:** Scan source-code repositories, CI/CD pipelines, infrastructure-as-code files, and container images for embedded credentials and LLM API keys. Rotate exposed secrets, invalidate associated tokens, and enable GitHub secret scanning and push protection.[1, 10]
- **Reduce internet-facing attack surfaces:** Inventory externally accessible applications, including partner portals, mobile backends, and legacy services. Restrict administrative interfaces through VPN or zero-trust access controls, disable anonymous file uploads, remediate SSRF vulnerabilities, enforce IMDSv2 on AWS workloads, and reject unsigned SAML assertions.[1, 10, 11]
- **Strengthen endpoint and identity protection:** Verify EDR coverage across Windows and Linux servers, require phishing-resistant MFA for user accounts, and apply appropriate Conditional Access and workload identity controls to limit unauthorized authentication and privilege escalation.[3, 9, 11]
- **Monitor suspicious outbound activity:** Investigate unexpected connections from enterprise servers to LLM API providers, including api.deepseek[.]com and open.bigmodel[.]cn, as well as webhook.site and other request-capture services. Correlate these connections with unusual process execution and outbound data transfers.[1, 11]

Structural

- **Automate identity containment:** Develop response playbooks that can disable compromised service principals, revoke application credentials, and invalidate associated access tokens. Require analyst review where necessary to prevent disruption to legitimate business services.[9]
- **Govern third-party application access:** Treat multi-tenant application permissions and OAuth consent grants as privileged access. Enforce least privilege, periodically recertify application permissions, remove unused integrations, and establish appropriate rate limits for business-critical APIs.[1, 10]
- **Establish visibility into authorized AI security testing:** Require internal red teams and approved security providers to register their use of ARTEX and other agentic penetration-testing frameworks with the SOC. Document approved infrastructure, testing windows, and expected activity so analysts can distinguish authorized assessments from malicious operations.[9]

© 2026 Critical Start. All rights reserved.

- **Improve detection of automated intrusions:** Develop behavioral detections for high-frequency reconnaissance, repeated authentication attempts, rapid permission enumeration, unexpected service principal modifications, and unusual access to cloud secrets and storage. Correlate activity across identity, endpoint, and cloud environments rather than relying exclusively on ARTEX-specific indicators.

Conclusion

ARTEX demonstrates how agentic AI is changing the speed and scale of cyber intrusions, allowing threat actors to automate reconnaissance, exploit existing weaknesses, and expand access across cloud and enterprise environments with limited human intervention. The risk extends beyond organizations directly targeted by attackers: compromised third-party application identities and excessive OAuth permissions can expose entire downstream customer environments. As ARTEX forks and similar frameworks become more widely available, organizations should expect continued adoption by both state-aligned and financially motivated actors. CRU recommends prioritizing visibility into application identities and third-party integrations, rapid credential revocation, and behavioral detection capable of identifying automated attack activity across cloud and on-premises environments. The ability to detect and contain compromised identities in minutes, rather than relying on traditional investigation timelines, will be critical to limiting the impact of agentic attacks.

For more threat reports, including OPTIMJACK and Java4Shell, visit Critical Start's Intel Hub. This advisory was written using the best intelligence available at the time and is subject to change as additional information becomes available.

What Critical Start is Doing

The CRITICALSTART® Cyber Research Unit (CRU) disseminated a brief of the threat as part of the 10-22-2026 and 11-05-2026 CTI Updates and is providing this Security Advisory as a follow up. CRU will continue to monitor ARTEX developments, as well as Agentic AI-accelerated attacks while working closely with the SOC and Security Engineering team to implement any relevant detections and increase vigilance. Critical Start customers are encouraged to reach out to support@criticalstart.com for any further enquiries.

Further Reading

[1] Federal Bureau of Investigation (Nov. 13, 2026). FLASH-20261113-001: Threat Overview of ARTEX Autonomous Penetration Testing Framework. https://www.ic3.gov/CSA/2026/261113.pdf

[2] Microsoft Threat Intelligence (Oct. 21, 2026). AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise. https://www.microsoft.com/en-us/security/blog/2026/10/21/storm-0940-uses-artex-agentic-ai-connectwise-supply-chain-compromise/

[3] CrowdStrike (Oct. 29, 2026). VAPOR PANDA Uses AI-Driven ARTEX to Scale a Global Espionage Campaign. https://www.crowdstrike.com/en-us/blog/vapor-panda-uses-artex-to-scale-global-campaign/

[4] CrowdStrike (Oct. 7, 2026). Unknown Threat Actor Uses AI-Driven ARTEX to Target South Korean Finance. https://www.crowdstrike.com/en-us/blog/unknown-threat-actor-uses-artex-to-target-south-korean-finance/

[5] Google Threat Intelligence Group and Mandiant (Oct. 30, 2026). ARTEX and Friends: Enter Agentic Offensive Frameworks. https://cloud.google.com/blog/topics/threat-intelligence/artex-and-friends-enter-agentic-offensive-frameworks

© 2026 Critical Start. All rights reserved.

[6] Check Point Research (Nov. 5, 2026). Hacking on Autopilot: Inside ARTEX Cybercrime Campaigns. https://research.checkpoint.com/2026/hacking-on-autopilot-inside-artex-cybercrime-campaigns/

[7] Microsoft Threat Intelligence (Nov. 6, 2026). ARTEX abuse expands as cybercriminals target financial sector. https://www.microsoft.com/en-us/security/blog/2026/11/06/artex-abuse-expands-as-cybercriminals-target-financial-sector/

[8] Censys (Nov. 5, 2026). ARTEX: Agentic Asbestos. https://censys.com/blog/artex-agentic-asbestos/

[9] Sophos X-Ops (2026). ARTEX: Attackers apply agentic AI. https://www.sophos.com/en-us/blog/artex-attackers-apply-agentic-ai

[10] SentinelOne (Nov. 13, 2026). ARTEX Unmasked: The Rise of AI-Powered ARTEX in Cybercrime and Espionage. https://www.sentinelone.com/blog/artex-unmasked-the-rise-of-ai-powered-artex-in-cybercrime-espionage/

[11] SOCRadar (Nov. 5, 2026). ARTEX AI Pentest Tool: What Security Teams Need to Know. https://socradar.io/blog/artex-ai-pentest-tool-what-security-teams-need-to-know/

[12] The Record (Nov. 13, 2026). FBI flash alert says Artex tool automating cyberattacks at machine speed. https://therecord.media/artex-ai-tool-fbi-flash-alert

Appendices

Appendix A: Indicators of Compromise

Indicator-based blocking alone is insufficient to contain ARTEX activity because much of its supporting infrastructure is cloud-hosted, shared, or transient. IP addresses should be treated as historical observations for threat hunting rather than permanent indicators of malicious activity. Security teams should corroborate them against current telemetry before blocking and investigate any historical connections to identified infrastructure.[1] Operators also use residential proxy networks to rotate their apparent source addresses. One proxy service, for example, showed 79 exit IP addresses across 100 sampled identities.[10] Consequently, blocking individual IP addresses may disrupt observed traffic without eliminating an attacker's access. Organizations should prioritize identifying and revoking compromised credentials, exposed secrets, unauthorized application identities, and associated access tokens.

Type	Indicator	Notes
IPv4	144.172[.]64.128	Storm-0940; sign-ins and Graph API activity through ConnectWise applications[10, 11]
IPv4	144.172[.]64.99	Storm-0940; password spray, October 10 to 14, 2026[2, 11]
IPv4	144.172[.]64.60	Storm-0940; failed user logons, October 18, 2026[10, 11]
IPv4	187.15[.]88.186 187.15[.]90.217 187.15[.]89.186	Storm-0940; VPN addresses used for Graph API access, October 15 to 16, 2026 (transient)[1, 10]
App ID	506e08f9-6e9a-4da4-942b-a099a03644ae	ConnectWise Cloud Backup; inventory and review[10]
App ID	24205314-2b89-4e07-a2a4-579e898a5640	ConnectWise, MS Teams application; inventory and review[10]

© 2026 Critical Start. All rights reserved.

Type	Indicator	Notes
App ID	0de14c77-b9b0-4137-9360-34878a929827	ConnectWise, MS Graph application; inventory and review[10]
Domain / IPv4	nginxmysql[.]com 45.153[.]126.223	VAPOR PANDA; BLOODALCHEMY command and control (C2)[3, 11]
Domain / IPv4	sqlserver-sqlite[.]com 89.46[.]234.50	VAPOR PANDA; BLOODALCHEMY C2[3, 11]
Domain / IPv4	www.ddduuuppp[.]com 192.121[.]47.20	VAPOR PANDA; BLOODALCHEMY C2[10, 11]
SHA-256	49fe70419fd19b81ca8d730a6ef7c1ec71b681d60b390761070d42686c7d2068	VAPOR PANDA; renamed GameHook binary (svchost.exe)[3, 11]
SHA-256	3fb5a9fb3e2cb1d6208966bd4e55d10282d19c7250d827efa560463d60a12ee4	VAPOR PANDA; BLOODALCHEMY loader graphics-hook-filter64.dll[3, 11]
SHA-256	364251838edfcf5ef1107fa692da3c5bd9a540482449f645992830134bf807ab	VAPOR PANDA; BLOODALCHEMY payload graphics-hook_config.bin[3, 11]
File path	C:\ProgramData\artx5v\	VAPOR PANDA staging directory[10, 11]
IPv4	38.244[.]50.120	South Korea campaign; ARTEX instance and open directory on port 18899[1, 4]
IPv4	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	South Korea campaign; proxy addresses (hunting leads only)[4]
Domain	xcai[.]pro	LLM API reseller used to reach DeepSeek[1, 4]
IPv4	199.19[.]95.66, 199.19[.]95.67, 199.19[.]95.68	Storm-3358; WISP C2 nodes[7, 10]
IPv4	91.92[.]240.238, 85.11[.]168.45, 178.16[.]54.129, 31.77[.]229.35	Storm-3358; ARTEX servers[7, 10]
IPv4	85.11[.]167.52	Storm-3358; C2 server[7, 10]
TLS cert	C7A884AFCD7F91408921890D9617BEF6BFAC3DAE992BC25BB9C48C0395A60842	ARTEX default TLS certificate (issuer organization ARTEX)[10]
Docker image	tomicdd/artex congdonglinux/artex	Backdoored ARTEX images[10]

Appendix B: Detection and Hunting

Behavioral indicators are more durable than infrastructure-based indicators because operators can readily rotate IP addresses, domains, and hosting environments. Security teams should prioritize identifying ARTEX-associated behaviors, including automated reconnaissance, credential harvesting, privilege escalation, and unauthorized cloud access. However, individual detections should be treated as investigative leads rather than confirmation of compromise, as authorized security testers also use ARTEX.[10] Correlate findings with user activity, authentication logs, application permissions, and changes to cloud identities to determine whether activity is malicious.

© 2026 Critical Start. All rights reserved.

Layer	Indicator
Cloud identity	Service principal sign-ins with scripted user agents; one address authenticating to three or more service principals with mixed success; one application requesting tokens across ARM, Key Vault, Storage, and Graph in a short window[9, 10]
Cloud and web	User agents including `artex-pentest/1.0`, `artex-audit/1.0`, `artex-readonly/1.0`, `ARTEX-worker/1.0`, or strings containing `ARTeX`[10]
Authentication	Password spray attempts at exact one-second intervals, waves of 20 to 60 minutes, and failed API calls retried shortly after with a different API version[10, 11]
Windows host	Files or paths beginning with `artx`, particularly under `C:\Windows\Temp`; tooling named `i<intentID>[variant]_<purpose>.<ext>` (for example `i12345_b2.ps1`); command markers such as `ARTX123`[9]
Web server	Files matching `artex-probe-*`; web shells such as `cl1217.jsp`, `f1219a.jsp`, `z1219.jsp`; unexpected *.aspx or *.jsp files; web server processes spawning shells[1, 11]
Container	Shell (`sh -c`) spawned by `/app/artex`[10]
Network	Server-originated requests to webhook.site or other request-capture services, and to LLM API endpoints or resellers[1, 11]

© 2026 Critical Start. All rights reserved.