VAPOR PANDA Uses AI-Driven ARTEX to Grow Global Espionage Campaign

By Counter Adversary Operations

Archived: 2026-11-03 02:00:56 UTC

CrowdStrike OverWatch identified a VAPOR PANDA campaign using ARTEX alongside the Chinese-developed DeepSeek large language model (LLM) and a second commercial LLM to conduct intrusions against multiple high-value targets in parallel.

Beginning on October 17, 2026, the adversary used ARTEX against more than 200 distinct organizations worldwide, including foreign ministries in Latin America, Europe, Asia, and the Middle East; semiconductor, defense, and think tank organizations in the U.S.; and telecommunications providers in multiple regions.

The VAPOR PANDA campaign detailed in this blog post is separate from the ARTEX activity recently reported by industry sources under the Storm-0940 designation.[1] CrowdStrike has previously documented ARTEX and similar agentic frameworks in use by financially motivated and unattributed actors; in this campaign, a targeted intrusion adversary used ARTEX to run parallel intrusions against multiple high-value targets. CrowdStrike Falcon telemetry also revealed details absent from public ARTEX reporting, including VAPOR PANDA's deployment of the *BLOODALCHEMY* backdoor and its command-and-control (C2) infrastructure.[2]

This campaign marks an escalation in China-nexus use of AI agents, from single intrusions to multiple simultaneous engagements directed by a human operator. The operation also shows the adversary building endpoint sensor evasion directly into their agentic workflow, including instructions to identify hosts without the CrowdStrike Falcon sensor.

Details

Targeting

Beginning October 17, 2026, CrowdStrike OverWatch observed exploitation activity from a single ARTEX controller directed against entities in at least 14 countries over an 11-day window. The controller, which CrowdStrike Intelligence has linked to VAPOR PANDA infrastructure, intermittently exposed adversary tooling, target files, and ARTEX agent transcripts. This indicates multiple engagements (referred to as “tasks” by ARTEX) against multiple targets.

VAPOR PANDA used ARTEX to conduct extensive intrusions at the foreign ministries of two Latin America-based countries, likely compromised a third, and staged target files for use against foreign ministries in Central America and Asia. The adversary also compromised a Middle East-based interior ministry. One engagement focused on a single European country, targeting its foreign ministry, presidency, and prime minister's office, as well as defense, justice, finance, and intelligence entities. A separate engagement targeted 19 U.S.-based

organizations in the pharmaceutical, semiconductor, defense, nuclear, and financial sectors. The adversary also compromised telecommunications providers in the U.S., Belgium, and the Philippines.

CrowdStrike Intelligence confirmed post-exploitation activity at nine organizations, where the adversary gained access, deployed *BLOODALCHEMY*, and conducted follow-on credential spraying, network scanning, and reconnaissance. The adversary also conducted password-spraying and credential-scanning activity against at least 12 additional organizations, including U.S.-based think tanks and government entities in Europe.

Inside the Operation

VAPOR PANDA ran ARTEX primarily with the Chinese-origin DeepSeek LLM, supplemented with a second, subscription-based commercial LLM that CrowdStrike Intelligence could not confirm. In a single intrusion, VAPOR PANDA's planning agents consumed on the order of 1.4 billion input tokens, indicating the scale of the automated activity and implying far higher consumption across the operation as a whole.

The campaign's reliance on open source components layered on top of commercial LLM subscriptions illustrates how adversaries can assemble scalable, multi-agent intrusions largely from publicly available resources.

ARTEX coordinates a main agent with subordinate planner and worker agents over a shared graph of facts, findings, and goals, advancing each engagement from initial access through host reconnaissance and lateral movement. The framework kept a human in the loop: Recovered transcripts showed an operator issuing instructions in Chinese — marked with the label 【人类指令】 ("human instruction") and timestamped in China Standard Time — and directing the agents to write web shells, extract domain administrator passwords, and inventory endpoint sensor coverage.

Where CrowdStrike has previously observed threat actors use AI agents to assist individual intrusions, VAPOR PANDA used ARTEX to run multiple engagements simultaneously, intervening only to steer the agents at decision points. The recovered transcripts and staged tooling show an adversary pairing AI-driven workflow with well-established tradecraft:

- At a U.S.-based pharmaceutical organization, the agent chained an unsigned Security Assertion Markup Language (SAML) assertion and an XML external entity (XXE) flaw to query the cloud Instance Metadata Service, obtaining a role capable of issuing administration commands to thousands of instances before escalating to recover an enterprise administrator account.
- At a second U.S.-based pharmaceutical organization, the agent deployed web shells, abused an unauthenticated orchestration service to obtain root access in a production container environment, read Local Administrator Password Solution (LAPS) passwords, and abused resource-based constrained delegation to impersonate a privileged account.
- Against a Greek government entity, the operator supplied the agent with a 2022 directory export from the organization's mail server containing authentication keys that had likely not been rotated since the export. The agent used those keys to forge an administrator token and read mailboxes, then used VPN configurations recovered from a user's sent mail, alongside a VPN client key the operator described as originating from a May 2024 email leak, to reach the internal network.

Across these engagements, the agents pursued multiple attack paths in parallel, spanning web applications, cloud services, container environments, Active Directory, and mail and VPN infrastructure. They shifted to alternative approaches when a path failed. This breadth likely enables a single operator to apply a wider range of techniques against more targets at once than would be feasible through manual operations.

VAPOR PANDA was also observed directing the agents to evade endpoint defenses. At a global pharmaceutical organization, the operator instructed the agents to “在域里寻找没有装猎鹰的windows机器，从tasklist和文件夹目录判断” (English translation: "find Windows machines in the domain without Falcon, judging by tasklist and folder listings"), using 猎鹰, which is People’s Republic of China (PRC) red-team slang for the CrowdStrike Falcon sensor. Incorporating sensor coverage checks directly into the agentic workflow allows the adversary to concentrate follow-on activity on unprotected hosts.

BLOODALCHEMY

At a Middle East-based interior ministry and a Latin America-based foreign ministry, VAPOR PANDA deployed *BLOODALCHEMY*, a backdoor first publicly documented in October 2023. The adversary used a legitimate GameHook binary renamed svchost.exe to load a malicious DLL, graphics-hook-filter64.dll, which decrypted the *BLOODALCHEMY* payload file graphics-hook_config.bin. The adversary staged copies of all three files in the exposed open directories. The implant carries an internal project path, D:\ProjectX\Merman_New\, that likely reflects the malware's in-house development name.

Across the observed activity, *BLOODALCHEMY* communicated with VAPOR PANDA infrastructure.These indicators do not appear in prior public ARTEX reporting and are also consolidated in the IOC table at the end of this post.

Assessment

This campaign is separate from the ARTEX activity reported by industry under the Storm-0940 designation, likely indicating the AI-based framework is quickly gaining popularity among multiple China-nexus adversaries.

Since 2021, VAPOR PANDA has primarily targeted government, think tank, energy, and extractive organizations across Europe, North America, Oceania, and the UK to develop long-term access. This campaign marks the adversary's first observed use of AI-driven intrusion tooling and a notable broadening of its targeting into the pharmaceutical, semiconductor, nuclear, and defense sectors, as well as foreign ministries, across new regions. Agentic tooling like ARTEX enables this scale of expansion at speed, allowing a small team to pursue a wider target set than manual operations would support, rather than a wholesale shift away from the adversary's established access-development mission. The operation’s reliance on DeepSeek and additional commercial LLMs, run through a purpose-built harness, likely reflects a deliberate effort to construct a resilient toolchain from multiple model providers rather than depend on any single service.

China-nexus adversaries will highly likely continue adopting agentic frameworks such as ARTEX over the next 12 months, pairing them with commercially available LLMs, and enabling small teams to quickly scale intrusions across many targets simultaneously.

Recommendations

- Rotate signing keys, preauthentication keys, and service credentials after any suspected data exposure, including keys contained in historical backups and directory exports.
- Enforce SAML signature validation, require IMDSv2, and restrict cloud roles' command-execution permissions to the minimum required instances.
- Hunt for ARTEX markers, including command-line parameters containing artx or artex strings, artx-prefixed cloud roles, art18- administration-command comments, and the C:\ProgramData\artx5v\ directory.
- Patch internet-facing appliances, particularly long-lived Linux-based servers, and ensure they run the latest available software and firmware.
- Deploy an advanced, real-time endpoint detection and response (EDR) agent, such as the CrowdStrike Falcon sensor, across the environment, including on servers considered inaccessible from the public internet.

Indicators of Compromise (IOCs)

Table 1. IOCs related to the information provided in this report

IOC	Description
nginxmysql[.]com	BLOODALCHEMY C2 domain
45.153.126[.]223	IP address of nginxmysql[.]com
sqlserver-sqlite[.]com	BLOODALCHEMY C2 domain
89.46.234[.]50	IP address of sqlserver-sqlite[.]com
www[.]ddduuuppp[.]com	BLOODALCHEMY C2 domain
192.121.47[.]20	IP address of www[.]ddduuuppp[.]com
49fe70419fd19b81ca8d730a6ef7c1ec71b681d60b390761070d42686c7d2068	GameHook binary renamed svchost.exe
3fb5a9fb3e2cb1d6208966bd4e55d10282d19c7250d827efa560463d60a12ee4	BLOODALCHEMY loader DLL (graphics-hook-filter64.dll)
364251838edfcf5ef1107fa692da3c5bd9a540482449f645992830134bf807ab	BLOODALCHEMY payload file (graphics-hook_config.bin)

MITRE ATT&CK

Table 2. MITRE ATT&CK tactics and techniques described in this report

Tactic	Technique	Observables
Resource Development	T1588.007 – Obtain Capabilities: Artificial Intelligence	Used the ARTEX agent framework to conduct intrusions
Initial Access	T1190 – Exploit Public-Facing Application	Exploited SAML, XXE
Initial Access	T1078 – Valid Accounts	Used stolen VPN configurations and leaked credentials
Initial Access	T1133 – External Remote Services	Connected to internal networks via VPN
Execution	T1651 – Cloud Administration Command	Issued commands to instances via a cloud systems manager
Persistence	T1505.003 – Server Software Component: Web Shell	Deployed JSP and PHP web shells
Privilege Escalation	T1134.001 – Access Token Manipulation: Token Impersonation/Theft	Abused resource-based constrained delegation to impersonate privileged users
Defense Evasion	T1574.001 – Hijack Execution Flow: DLL	Executed the BLOODALCHEMY loader DLL via a legitimate, renamed binary
Defense Evasion	T1036.005 – Masquerading: Match Legitimate Resource Name or Location	Renamed a legitimate GameHook binary as svchost.exe
Credential Access	T1552.005 – Unsecured Credentials: Cloud Instance Metadata API	Queried the metadata service via XXE to obtain cloud role credentials
Credential Access	T1555 – Credentials from Password Stores	Read LAPS passwords
Credential Access	T1606.002 – Forge Web Credentials: SAML Tokens	Submitted unsigned SAML assertions
Credential Access	T1606 – Forge Web Credentials	Forged a mail server administrator token with unrotated keys
Credential Access	T1110.003 – Brute Force: Password Spraying	Password-sprayed government, think tank, and corporate accounts
Discovery	T1518.001 – Software Discovery: Security Software Discovery	Inventoried endpoint sensor coverage

Collection	T1114.002 – Email Collection: Remote Email Collection	Read mailboxes on a compromised mail server
Command and Control	T1090.003 – Proxy: Multi-hop Proxy	Routed traffic through Tor, SOCKS proxies, and a compromised relay
Command and Control	T1572 – Protocol Tunneling	Tunneled traffic via Chisel and SoftEther

[1] https://www.microsoft.com/en-us/security/blog/2026/10/21/storm-0940-uses-artex-agentic-ai-connectwise-supply-chain-compromise/

[2] https://www.elastic.co/security-labs/threat-command/disclosing-the-bloodalchemy-backdoor

---

Source: https://www.crowdstrike.com/en-us/blog/vapor-panda-uses-artex-to-scale-global-campaign/