MORPHISEC

THREAT ADVISORY

ARTEX: Agentic AI Hits the Endpoint

What happens when autonomous attackers reach your hosts

November 12, 2026 | TLP:GREEN | Morphisec Threat Labs

Autonomous AI agents are accelerating every phase of the attack chain. But the chain has a bottleneck: to persist, steal credentials, or move laterally, the attacker must execute code on a host. This advisory examines the three threat clusters using the ARTEX framework in Q4 2026, maps where their operations touch endpoints, and assesses Morphisec prevention at each host-level stage.

200+
ORGANIZATIONS TARGETED

3
THREAT CLUSTERS

60s
CREDENTIAL PIVOT TIME

1,500
COMMANDS PER HOUR

THE THREAT

ARTEX: From Pentest Tool to Threat Multiplier

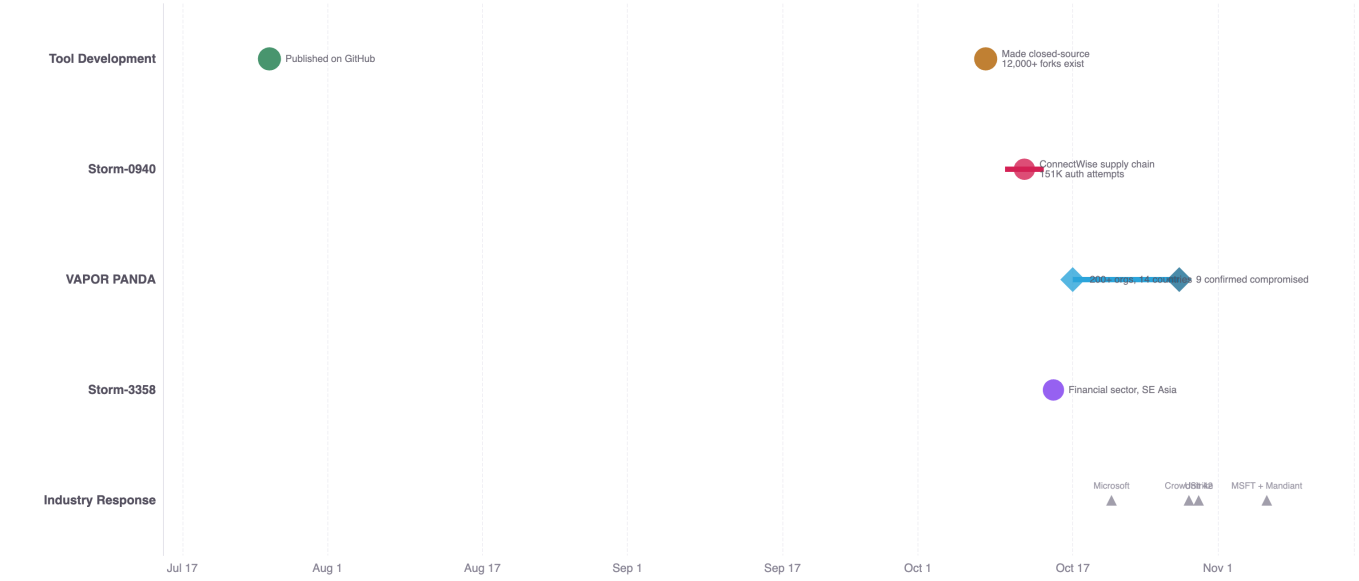
ARTEX is an open-source, LLM-driven penetration testing framework published on GitHub in July 2026. It coordinates AI planner agents with worker agents that run conventional tools (Nmap, Nuclei, sqlmap, Playwright). The planner chooses the next step, workers execute, and results feed into a shared graph. When one path fails, the system pivots to another within seconds. By October 2026, three unrelated threat clusters had adopted it.

What makes ARTEX operationally significant is the speed. CrowdStrike observed 1,500 commands per hour from a single ARTEX controller, compared to 40 per hour in historical hands-on-keyboard operations. Credential pivots happened within 60 seconds of discovery. In one intrusion, the planning agents consumed 1.4 billion input tokens. The original repository was made closed-source on October 8 after 12,000+ forks. The code continues to proliferate.

ARTEX agents handle reconnaissance, API enumeration, and identity attacks in the cloud. But the attack chain has a bottleneck: to persist, harvest credentials, or move laterally, the attacker must land on a host. That host-level activity is where endpoint prevention controls engage.

WHO IS USING ARTEX

ACTOR	ORIGIN	TARGETS	HOST-LEVEL ACTIVITY	SOURCE
VAPOR PANDA	China	200+ orgs across 14 countries: foreign ministries, defense, semiconductors, telecoms	Deployed BLOODALCHEMY backdoor, DCSync, web shells, Chisel/SoftEther tunnels	CrowdStrike, Oct 29
Storm-0940	China	ConnectWise customers, 2,250 M365 tenants	Password spray (151K attempts, 2,900/min), credential harvesting	Microsoft, Oct 21
Storm-3358	Cambodia	Banks and financial institutions, SE Asia + global	Deployed WISP, vshell, AdaptixC2 RATs via netcat and curl	Microsoft, Nov 6



Campaign timeline: published July 26, three threat clusters by October, closed-source October 8 after 12,000+ forks. Activity is ongoing.

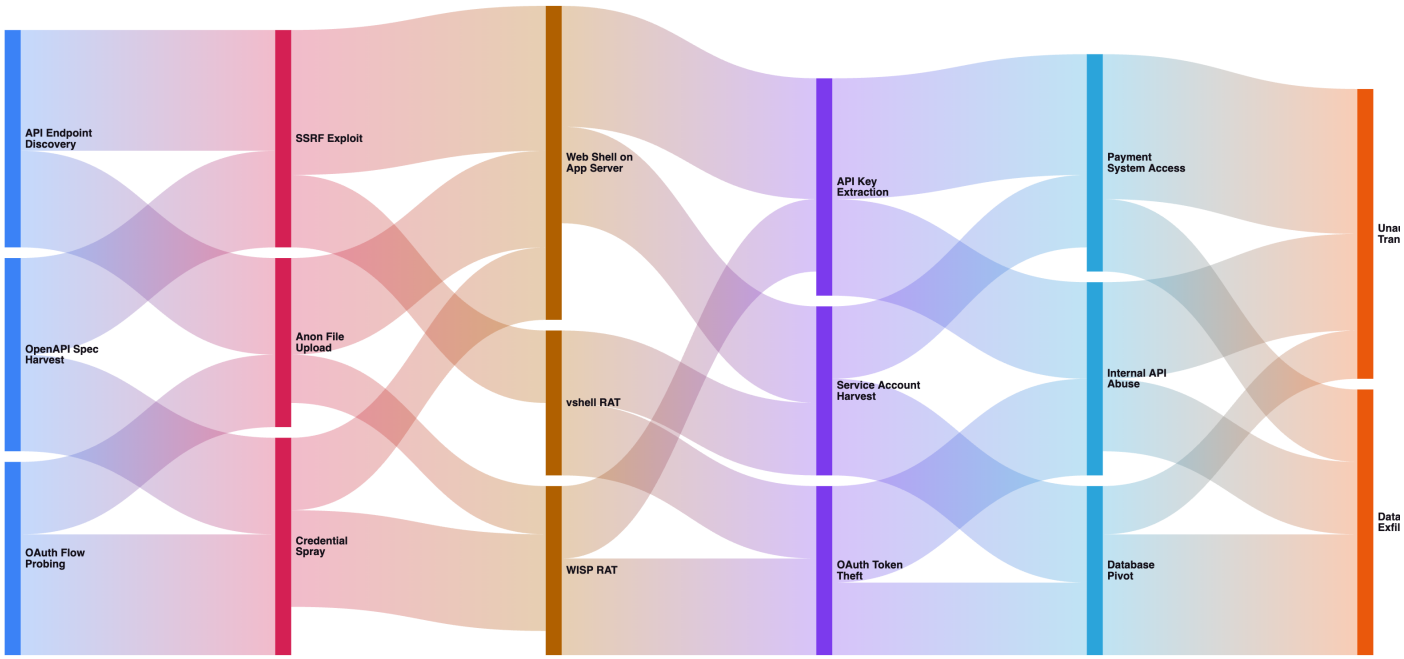
WHEN ARTEX TOUCHES THE HOST

The Four Stages That Matter for Endpoint Security

ARTEX agents do extensive work before touching a host: scanning, API enumeration, OAuth probing, credential spraying. Those stages are primarily network and identity events. Endpoint prevention engages when the attacker lands on a host to stay or to move. Across the three campaigns observed, four host-level stages appeared consistently.

STAGE	WHAT ARTEX AGENTS DO	OBSERVED TOOLS	MORPHISEC
Implant deployment	Download and execute RATs, backdoors, and web shells on compromised servers. Delivery via curl, netcat, or direct file write after exploiting a vulnerability.	WISP, vshell, AdaptixC2, BLOODALCHEMY, custom web shells	**Prevents**
Credential harvesting	Extract credentials from memory, registry, browser stores, and Active Directory. Used to pivot to payment systems, databases, and internal APIs.	DCSync, LAPS password reads, AD reconnaissance, Mimikatz-style tools	**Prevents**
Lateral movement	Use harvested credentials to reach additional hosts. Move from compromised app servers toward databases, domain controllers, and high-value targets.	PsExec, Chisel tunnels, SoftEther VPN, resource-based constrained delegation	**Prevents**
Destructive payloads	Deploy ransomware, wipe shadow copies, or tamper with security controls. Used by financially motivated clusters (Storm-3358) after data staging.	Ransomware loaders, shadow copy deletion, security service tampering	**Prevents**

The Sankey diagram below traces the full attack flow from reconnaissance tools through exfiltration, showing where host-level events occur. The right half of the diagram, from implant deployment onward, is where Morphisec prevention applies.



MORPHISEC PREVENTION

What Gets Stopped at the Host

THREAT	OUTCOME	STATUS
RAT deployment (WISP, vshell, AdaptixC2, BLOODALCHEMY)	Prevented on Windows. Prevented on Linux via trust rules (unauthorized binaries blocked at execution).	**Prevent**
Web shell execution on application servers	On Windows, payload delivered by the web shell is blocked in memory. On Linux, unauthorized scripts are blocked by trust rules if outside the approved baseline.	**Prevent**
Credential theft from OS stores (LSASS, SAM, browser vaults)	Decoy credentials trigger alerts when accessed. Credential harvesting tools encounter fake targets before reaching real stores.	**Prevent**
DCSync and Active Directory credential extraction	Decoy domain controller objects detect DCSync attempts. Alerts fire before the attacker obtains usable domain credentials.	**Detect**
Lateral movement via PsExec, SMB, remote admin tools	PsExec-style tools blocked. Remote encryption attempts (ransomware over SMB) prevented.	**Prevent**
Fileless and living-off-the-land execution	Prevented regardless of whether the payload arrives as a file, a PowerShell download cradle, or a reflective injection.	**Prevent**
Ransomware and destructive payloads	Shadow copy protection prevents deletion. Encryption decoys trigger early termination. Key material captured from process memory during encryption.	**Prevent**
SIEM integration	Events for every blocked attack, decoy trigger, and policy violation. Forwarding to Splunk, Sentinel, QRadar, and syslog-compatible SIEMs.	**Detect**



DETECTION GUIDANCE

Behavioral Signals of ARTEX Operations

ARTEX is forked and modified constantly. Static indicators lose value fast. The behavioral patterns below are durable across all observed ARTEX campaigns.

1. Inference-Paced Request Bursts

Bursts of 1 to 14 requests at ~1s intervals, then 25 to 31 second pauses. This cadence on auth, registration, or API endpoints is the fingerprint of an LLM-backed agent waiting on inference.

2. Canary Usernames in Identity Provider Logs

Registration or auth attempts using evil[.]com, test@example[.]com, or deadbeef hex patterns. These precede real credential testing.

3. OpenAPI Spec Downloads Followed by Structured Probing

Automated downloads of API specifications followed immediately by targeted requests against every discovered route.

4. First-Use Cloud Key from New IP

A service account or API key used from a new IP address, especially if followed by role escalation or data access within minutes.

5. Write-Then-Delete Storage Tests

ARTEX agents test permissions before staging exfiltration. A write followed by an immediate delete on any object store is a strong signal.

6. RAT Downloads to Server Directories

Curl or netcat downloading executables to web-accessible directories. Match against known hashes for WISP, vshell, and AdaptixC2.

WHAT ENDPOINT SECURITY DOES NOT COVER

ARTEX operates across cloud, identity, and network layers before reaching a host. Cloud-only attack chains (stolen OAuth secrets used directly against cloud APIs) never touch an endpoint and generate no endpoint telemetry. API keys stored as environment variables or config files on disk are not protected by credential decoys. These paths require complementary controls: secrets management, cloud posture management, and identity threat detection. Morphisec is the host-level layer, not a replacement for those.

RECOMMENDATIONS

Prioritized Actions

1. Deploy Morphisec on all application servers and workloads that process financial transactions. On Windows, in-memory attacks are stopped before execution. On Linux, configure trust rules to block unauthorized binaries and scripts.
2. Deploy credential decoys across application servers. Autonomous AI agents cannot distinguish decoy credentials from real ones and will waste compute pursuing false leads. This turns speed against the attacker.
3. Move API keys and service credentials into a secrets manager (HashiCorp Vault, AWS Secrets Manager, Azure Key Vault). This removes the credential theft path that endpoint decoys do not cover.
4. Remove OpenAPI specifications, source maps, and configuration files from internet-facing applications. These are target lists for ARTEX agents.
5. Require authentication on every write endpoint, including development and staging APIs.
6. Rate-limit and bot-protect token issuance, password reset, and registration endpoints.
7. Implement transaction velocity monitoring on disbursement APIs. Alert when legitimate credentials issue requests at rates inconsistent with normal application behavior.
8. Hunt for agentic behavior patterns (inference-paced bursts, canary inputs, write-then-read verification) rather than static indicators. Behavior is durable. Indicators are not.

MITRE ATT&CK

TECHNIQUE	ID	TECHNIQUE	ID
Exploit Public-Facing App	T1190	Password Spraying	T1110.003
Valid Accounts	T1078	Web Shell	T1505.003
OS Credential Dumping	T1003	PowerShell Execution	T1059.001
SMB / Admin Shares	T1021.002	Protocol Tunneling	T1572
Security Software Discovery	T1518.001	Exfiltration Over C2	T1041

SOURCES

Unit 42, "ARTEX and AI-Powered Red Teaming Tools," November 2026. unit42.paloaltonetworks.com
CrowdStrike, "VAPOR PANDA Uses AI-Driven ARTEX to Scale a Global Espionage Campaign," October 29, 2026. crowdstrike.com
Microsoft, "ARTEX abuse expands as cybercriminals target financial sector," November 6, 2026. microsoft.com