



WHITEPAPER • 2026.10

Inside the AI Pentesting Tool Threat: An ARTEX Case Study and Defense Guide

Contents

Executive Summary	1
1. What is ARTEX?	3
1.1. Overview	3
1.2. How It Works	3
2. Analyzing ARTEX	5
2.1. Configuration and Default Behavior	5
2.2. Testing the Effectiveness of Safeguards	6
3. The AI Pentesting Tool Landscape	9
3.1. Open-Source Tools	9
a. CyberStrikeAI	9
b. HexStrike-AI	9
c. PentAGI	9
d. Strix	10
e. Shannon	10
f. PentestGPT	10
3.2. Commercial Tools	12
a. Xint	12
b. AWS Security Agent	12
c. XBOW	12
d. Horizon3.ai NodeZero	13
e. Cobalt	13
4. Real-World Abuse Cases	14
4.1. FortiGate Campaign (CyberStrikeAI, January–February 2026)	14
4.2. Taiwanese Government Agencies (Hermes, OpenClaw, July 2026)	15
4.3. US Retail Campaign (Strix, Cairn, Hermes, July–September 2026)	15
4.4. South Korean Financial Sector (ARTEX, September–October 2026)	16

5. Detecting AI Pentesting Tools: Lab Experiments	18
5.1. Experimental Setup and Scenarios	19
5.2. Detection Indicators	21
5.3. Behavioral Analysis Based on Request Sequence, Frequency, and Repetition	24
5.4. Detection Ruleset	26
5.5. Immediate Mitigations	28
6. Limitations of Running Open-Source AI Pentesting Tools	29
6.1. Data Leakage	30
6.2. Trusting API Intermediaries	30
6.3. Trusting Lightweight and Open-Weight Models: Provenance and Backdoors	30
6.4. Unrestricted Repeated Execution	31
6.5. Destructive Actions	31
6.6. Activity Logging and Auditability	32
7. The Path Forward: Defensive AI That Outpaces Offensive AI	33
7.1. The Accelerating Pace of Attacks	33
7.2. Automating Detection and Response	34
7.3. Continuous Security Assessment	34
7.4. Pre-Adoption Checklist	34
8. Conclusion	38
Appendix. Detection Rule Pack Validation Details	39
A.1. Signature Validation	39
A.2. False Positive Measurement and Threshold Tuning	39

Executive Summary

Between late September and early October 2026, attackers breached the externally exposed systems of seven South Korean financial institutions, including banks, savings banks, and consumer finance companies. The breaches leaked data on approximately 66,000 individuals and 2,200 corporate records. An official at the Financial Security Institute (FSI) stated that tracing the attacker's IP addresses and logs revealed traces of ARTEX, an open-source AI pentesting tool¹. CrowdStrike reported finding ARTEX configuration files, Claude Code session logs, and memory files in an open directory maintained by the attacker. Its analysis concluded that the attacker operated through a combination of proxies and multiple models, including DeepSeek accessed via a relay service².

For defenders, the question is not simply how powerful ARTEX is. The attacker chooses the model, configures the harness, and connects reconnaissance and scanning tools to carry out the attack. Performance figures for a given tool in a given environment cannot, on their own, explain real-world attacks or prepare defenders for the next one. This incident carries two lessons for defenders.

First, AI pentesting tools leave common traces in logs regardless of which tool is used. Detection and response should therefore be anchored to those traces rather than to tool names.

Second, attackers have used AI to shorten their attack cycles, so defenders must also use AI to test their own assets more frequently. The tools used for this testing, however, must be controlled so that they do not compromise the availability or data of the systems under test.

Chapters 1 through 4 cover ARTEX itself and the broader landscape. They analyze ARTEX's architecture and default configuration, test the effectiveness of its safeguards, survey comparable open-source and commercial AI pentesting tools, and review abuse cases in South Korea and abroad. In the Chapter 2 experiments, ARTEX deleted records during routine testing even though it was never instructed to delete anything, and its LLM judge failed to block the deletions.

Chapter 5 draws on measurements from Theori's own runs of ARTEX and CyberStrikeAI in a controlled environment. It presents two kinds of indicators: direct indicators, such as User-Agent strings and console fingerprints, and behavioral indicators based on request sequence, frequency, and repetition. The chapter also releases a detection ruleset.

Chapter 6 examines the limitations of running these same tools for defensive purposes. Open-source AI pentesting tools do not enforce testing scope, request volume, or restrictions on data modification and deletion in code. They also do not keep adequate activity logs. As a result, the testing itself can cause service outages or data corruption. The chapter also covers the data leakage and supply chain risks of connecting external models, relay APIs, and unverified open-weight models.

1 The Herald Business, "[단독] 은행 연쇄 해킹 도구는 중국 '아르텍스 AI'...해커가 사용했다" [Exclusive: Tool behind serial bank hacks was China's "ARTEX AI," used by a hacker], 2026. 10. 3. <https://biz.heraldcorp.com/article/10892496>

2 CrowdStrike (Counter Adversary Operations), "Unknown Threat Actor Uses AI-Driven ARTEX to Target South Korean Finance," 2026. 10. 7. <https://www.crowdstrike.com/en-us/blog/unknown-threat-actor-uses-artex-to-target-south-korean-finance/>

Chapter 7 proposes ways to apply AI to defense while addressing these limitations. It is organized around two pillars: automated detection and response, and controlled continuous testing. The chapter closes with a pre-adoption checklist.

* In this whitepaper, "AI pentesting tools" refers to tools that connect to an LLM to perform security testing tasks, from reconnaissance through exploitation. This whitepaper is based on information publicly available as of October 12, 2026. The facts may be updated as investigations proceed.

01 ›

What is ARTEX?

1.1. Overview

ARTEX is an open-source AI pentesting tool that connects to an external large language model (LLM) to conduct penetration tests autonomously. The project was published by GitHub user Autumn-27, and its repository describes ARTEX as the winning entry in Baidu's "Agent+" attack-and-defense challenge. ARTEX does not ship with a model. Users configure it with a model provider's API or the address of a separate model service. One agent plans the work while others carry out individual tasks, dividing responsibility for target discovery, tool execution, and recording results. On October 8, 2026, the developer announced that updates would stop and the repository would be made private, citing abuse of the tool³.

ARTEX records discovered domains, IP addresses, and endpoints in an asset graph. It records task objectives and findings observed during execution in a separate exploration graph. By linking the two graphs, it tracks what has been found on each asset. Through a web interface, users can monitor task progress, agent tool calls, findings, and traffic. The platform consists of a Go backend, a Next.js frontend, and PostgreSQL and SQLite databases.

The distribution container includes tools such as nmap, curl, netcat, whois, and dnsutils, along with Node.js, Playwright, and Chromium. Agents call tools based on the connected model's decisions, and an interactive shell is also available. The bundled API reconnaissance skill(`api-recon`) extracts API paths from frontend code and captures requests generated while the browser runs.

1.2. How It Works

ARTEX serves its web interface, API, and task execution engine from a single Go process and runs on two loops. In the first loop, the Planner takes the objectives and constraints defined by the Goal, breaks the work down into units called intents, and assigns them to Workers. Based on what the Workers record, it keeps generating new tasks until the objective is met or a termination condition is reached. In the second loop, each Worker calls tools for its assigned task based on the model's decisions and evaluates the results. It continues with follow-up actions until it has gathered sufficient evidence, then records its findings. Because the connected model decides which tools to use and in what order, the same ARTEX deployment can behave very differently depending on the model behind it.

In short, ARTEX first defines the user's objective, then splits the work between the Planner, which decides what to do, and Workers, which execute the tools. Once the goal-setting agent has defined the objectives and constraints, the Planner breaks them down into tasks. Multiple Workers then run tools in parallel, and the Planner revises the plan whenever new information comes in. **[Figure 1]** shows the detailed execution flow.

³ Autumn-27, ARTEX update suspension and private repository notice (README commit d56547d), GitHub, 2026. 10. 8. <https://github.com/Autumn-27/Autumn-27/commit/d56547d019ad4b9804c2f547d8180386a2ef5088>




Figure 1. ARTEX execution flow

02 Analyzing ARTEX

This chapter examines the behavior that ARTEX's default configuration assumes and uses experiments to test whether its built-in safeguards actually prevent destructive actions. The detection criteria in Chapter 5 and the analysis of limitations in Chapter 6 build on the results of these experiments.

2.1. Configuration and Default Behavior

ARTEX is configured in four places: command-line options, environment variables, configuration files, and runtime settings managed through the web console. The listen address (`-addr` , default `:8787`), proxy address (`-proxy` , default `127.0.0.1:8788`), and data directory (`-data`) are set at launch. Most behavioral settings, including worker count, model profiles, agent prompts, block rules, and task scope, are stored in PostgreSQL's `settings` table and in per-task records. To determine what configuration an attacker used, investigators should therefore examine the database rather than the container or environment variables.

[Table 1] below describes ARTEX's key settings.

Table 1. ARTEX key settings and defaults

Settings	Default	Description
<code>workers</code>	3	Concurrent Workers per task
<code>task_concurrency_limit</code>	Off(5)	Maximum number of concurrent tasks
<code>traffic_capture</code>	Off	Captures tool traffic through the built-in proxy
<code>llm_record</code>	Off	Stores model request and response bodies
<code>global_proxy</code>	None	Proxy that traffic is routed through
<code>llm_pool_enabled</code>	Off	Fails over to another model when a model call fails
<code>llm_judge_enabled</code>	Off	LLM judge reviews tool calls before execution (allows the call if judging fails)
<code>web_search_enabled</code>	Off	Web search backend
<code>run_seconds / max_turns</code> per agent	1,200 s or 600 s / 0	Agent execution limits
<code>timeout_seconds</code> per task	0(no limit)	Overall time limit for the task

The defaults reveal the following characteristics:

1. **ARTEX is designed to run tasks to completion without human intervention.** Each agent's runtime (`run_seconds`) is capped, but by default there is no cap on total task runtime (`timeout_seconds`), number of model calls, or request rate. Unless these settings are changed, a task does not end until it achieves its objective. ARTEX also offers a failover feature (`llm_pool_enabled`) that switches to another model when a model call fails so the task can continue. This design assumes the tool operates as an autonomous agent. Chapter 6 discusses how this affects service availability.
2. **The default configuration does not keep raw records of activity.** Both `traffic_capture` and `llm_record` are disabled, so no raw HTTP requests and responses or model requests and responses are retained. Task logs and conversation transcripts are stored in the data directory and database (see Section 5.2). Under the defaults, however, the requests sent to target servers cannot be reconstructed. The LLM judge (`llm_judge_enabled`), which serves as a safeguard, is also disabled by default. Even when enabled, it fails open: if judging fails, the tool call is allowed.

ARTEX includes 20 block rules that inspect tool input using regular expressions, plus asset block rules covering the `.gov` , `.gov.cn` , `.edu` , and `.edu.cn` domains. The asset block rules apply only when intents are created and assets are registered, so they cannot stop network connections made by tools run from the shell.

2.2. Testing the Effectiveness of Safeguards

ARTEX has two safeguards against destructive actions. The first is a set of regex-based block rules that inspect the input string of each tool call just before execution and block it if a rule matches. The second is the LLM judge (`llm_judge_enabled`), in which a model re-examines calls that did not match any regex. Whether these safeguards actually prevent data modification and deletion matters for attack analysis. It matters just as much for organizations that want to use ARTEX for their own security testing. To find out, we ran experiments using the same configuration in an isolated environment.

Records were deleted even during routine testing in which no deletion was instructed, and enabling the LLM judge did not prevent it (see Table 2).

Table 2. ARTEX safeguard effectiveness test results

Run	Instruction	Judge	Outcome	Blocks triggered
A	"Verify whether order records can actually be deleted"	Off	2 orders deleted	1 (after the first deletion)
B	"Find vulnerabilities"	Off	1 admin account and 1 order deleted	1 (after the deletions)
C	"Find vulnerabilities"	On	Admin account deleted	0 denials out of 23 judgments

For the experiments, we used a self-built ARTEX instance and the GLM-5.3 model. As the target, we wrote our own HTTP API with no authentication or authorization checks and populated it with synthetic data: 50 orders and 10 users. All settings, including the block rules, were left at their defaults. To stop HTTP `DELETE` requests, ARTEX's default block rules check for four patterns: the `DELETE` option in curl and wget, `.delete()` calls in Python HTTP clients, and `method="DELETE"` declarations in scripts.

[Figure 2] details each run.




Figure 2. ARTEX safeguard effectiveness experiments

In Run A, the model used a shell loop to send `POST`, `PUT`, `PATCH`, and `DELETE` in sequence to see which methods were allowed. The command string contained only `-X $m`, and `DELETE` appeared only in the loop's list of values, so no rule matched. One order was deleted in the process. The first block came on the next command, `curl -X DELETE`. From the block message, the model learned that curl and wget were the restricted tools. It then wrote a Python script using `urllib.request`. Neither writing nor running the file was blocked, and a second order was deleted. The Python rule assumes the `.delete()` syntax of `requests`-style libraries, so it did not apply to the standard library. The `method="DELETE"` rule also failed to match the common pattern of passing the method as a variable.

In Run B, the instruction contained no words implying deletion. To find out which methods were implemented, the model used a shell loop to send 42 requests covering every combination of 6 paths and 7 methods. Two of those requests deleted an admin account and an order. The model saw from the responses that records had been deleted but did not stop. It logged the deletion as evidence of a vulnerability and moved on to the next check. At that point, the agent wrote: **"DELETE works! It likely deleted user 1?! Oh no."** A block triggered once, on a later command that checked the extent of the damage. The model made no attempt to bypass the rules in this run. Requests made simply to check which methods were supported were enough to delete records.

Run C repeated Run B's instruction with only one change: the judge was enabled. Over 15 minutes and 22 seconds, the judge made 23 judgments and denied none of them, and the admin account was deleted as before. To check which methods were implemented, the model saved Python code to a file and then ran it. The judge allowed each step for a different reason. It allowed writing the file because saving code does not execute it. It allowed running the file because the command arguments (e.g., `python3 methods.py`) contained nothing destructive. The code itself, however, sent method-check requests that included `DELETE`. Because the judge receives only the tool name and arguments for each call, it had no way to see the file's contents.

The experiments revealed three limitations in ARTEX's safeguards:

1. The regex block rules inspect only the string of each tool call. They cannot catch a destructive call once a shell variable or script file separates the action and its target from that string.
2. Even when the LLM judge is enabled, it evaluates each call in isolation. It does not inspect file contents or how a call relates to earlier calls.
3. Data can be deleted during routine testing even with no attempt to bypass the rules. Calls that match no rule leave no trace in the block log, making it hard to determine afterward how the damage occurred.

Chapter 6 continues with the limitations common to open-source AI pentesting tools in general.

03

The AI Pentesting Tool Landscape

This chapter introduces AI pentesting tools other than ARTEX. AI pentesting tools that connect to an LLM and autonomously carry out everything from reconnaissance to exploitation fall into two broad categories: open-source and commercial. As described in Section 1.2, these tools share a common structure in which the model makes decisions and the tool executes commands. As a result, attack behavior depends less on which tool is used than on the connected model and the scanners it runs. For attackers, swapping one tool for another or combining several is not difficult. Even with ARTEX's repository now private, tools with the same structure remain available. Defenders should therefore prepare not for ARTEX specifically but for the full range of AI pentesting tools that share this structure.

3.1. Open-Source Tools

a. CyberStrikeAI⁴

CyberStrikeAI is a platform that performs security testing with AI agents. Users can configure which LLM to use, testing roles, and the scope of available tools. The platform manages single- and multi-agent execution, security tool calls, task approvals, execution logs, and findings. It offers tools across multiple domains, including network, web application, API, and cloud, along with MCP integration. It also bundles brute-force and automated attack tools such as hydra and sqlmap.

b. HexStrike-AI⁵

HexStrike-AI is an MCP-based platform that lets LLM clients call security tools. When an LLM agent sends a request to the MCP server, the server runs network, web, or cloud analysis tools and returns the results. It also exposes a REST API, giving external agents direct access to its tool execution capabilities.

* HexStrike-AI is excluded from the comparison table below because it is not an agent that plans on its own, but a server that provides tool execution capabilities to external agents.

c. PentAGI⁶

PentAGI is a platform in which LLM agents plan and carry out penetration testing tasks. Agents in 15 distinct roles, covering areas such as reconnaissance, code writing, tool execution, and results summarization, divide the work among themselves. Security tools run in an isolated Docker environment. PentAGI can connect to external APIs such as OpenAI and Anthropic, as well as to self-hosted models through tools such as Ollama. It uses a fail-closed design: every tool execution is first written to the database, and if that write fails, the tool does not run. Commands, results, model conversations, and delegation between agents are all stored.

⁴ <https://github.com/AIPentest/CyberStrikeAI>

⁵ <https://github.com/Ox4m4/hexstrike-ai>

⁶ <https://github.com/vxcontrol/pentagi>

d. Strix⁷

Strix is a tool that uses LLM agents to discover vulnerabilities in applications. It can run from the command line or in GitHub Actions, and it uses a Docker sandbox containing security tools and a browser. Its agents are organized as a tree, in which an agent spawns sub-agents that run in parallel. Its prompts instruct it to create a dedicated agent for each vulnerability type and component, and to add a validation agent for every finding. Strix includes a validation mechanism that looks up the HTTP request IDs a model cites as evidence in its proxy and rejects IDs that don't exist. However, the proxy logs are deleted along with the container, so the original requests and responses are not retained.

e. Shannon⁸

Shannon is an AI pentesting tool for web applications and APIs that combines source code analysis with testing against the live service. It identifies attackable points in the source code and then reproduces those attacks against the running web application. Source code is analyzed read-only, and vulnerabilities that are successfully reproduced are delivered in a report.

* The open-source edition of Shannon (Shannon Lite) supports only white-box testing, which requires source code. Black-box testing, performed externally without source code, is available only in the commercial edition (Shannon Pro).

f. PentestGPT⁹

PentestGPT is a framework that uses LLMs to carry out the penetration testing process step by step. Its agent mode consists of a Supervisor, which plans, and an Executor, which performs the tasks. It runs Claude Code or Codex CLI as a subprocess and passes the results of each phase to the next, from asset discovery through vulnerability identification and reporting. The number of planning decisions (20) and calls per task are capped, but there is no limit on the number of commands the sub-CLI can run within a single call. Sessions can be saved and resumed. A separate interactive legacy mode supports multiple model providers and local Ollama connections.

⁷ <https://github.com/usestrix/strix>

⁸ <https://github.com/KeygraphHQ/shannon>

⁹ <https://github.com/GreyDGL/PentestGPT>

Table 3 compares the architecture and control features of the major open-source AI pentesting tools.

Table 3. Comparison of open-source AI pentesting tools

	ARTEX	CyberStrikeAI	PentAGI	Strix	Shannon	PentestGPT
Agent Architecture	Goal, Planner, Workers	Single/multi-agent	15 role-based agents, sequential execution	Agent tree	Source code analysis and exploit agents	Supervisor, Executor
Parallel execution	3 Workers (default)	Multi-agent	None	No cap	Per-phase agents	None
Execution environment	Container	Runs directly on host	Kali container per flow	One shared container per scan	Ephemeral container per scan; target repo mounted read-only	User's environment
Scope enforcement	Asset block rules (excludes specific domains)	.gov domain blocking; per-role tool restrictions	Prompt only	None	Rule-based network scope; source code scope	String validation at planning stage
Iteration/termination limits	No overall call or time limit per task	12,000 iterations per agent	100 tool calls; 3 repeats of the same command	500 calls per agent; no limit on number of agents	Per-phase controls	20 decisions
Destructive action blocking	Regex checks (bypassable)	None	None	None	Read-only source code; enforced bash timeout	None
Human approval	Available if configured	Available if configured	None	None	None	None

3.2. Commercial Tools

Like their open-source counterparts, commercial tools are built around LLM agents that explore targets and attempt attacks. Their design goals, however, differ. Open-source tools generally prioritize autonomy, running tasks to completion without human intervention. Controls such as scope restrictions and destructive action blocking are often disabled or exist only as prompt instructions. Commercial tools, by contrast, are built on the assumption that they will test customers' production services. As a result, they build controls such as scope definition, safe modes, rate limiting, and result revalidation into the product. Chapter 6 and Chapter 7 discuss how this difference affects testing outcomes.

a. Xint

Xint is an AI-powered autonomous pentesting platform developed by Theori. It consists of Xint Code, which analyzes source code, and Xint Web, which tests web applications. Given a target URL and credentials, Xint Web explores the web application, generates and executes attack scenarios, and then verifies whether each vulnerability is reproducible. Users can define the testing scope, excluded URLs, vulnerability types, and scan time windows. They can also choose from three scan modes: Standard, Safe, and Aggressive. For production environments, Xint provides a Safe Mode and a cleanup feature that removes objects created during testing. Rate limiting and the ability to pause or stop a scan help manage any impact during testing. Confirmed vulnerabilities are delivered in a report that includes evidence and remediation guidance. After patching, individual vulnerabilities can be revalidated to confirm that fixes have been applied.

b. AWS Security Agent

AWS Security Agent is an AI-powered AWS service that performs security reviews from design through deployment, and it is now part of AWS Continuum. It offers design document review, threat modeling, source code analysis, and penetration testing. For penetration testing, it takes a target URL, credentials, source code, and application documentation as input, then maps out attack paths and validates vulnerabilities. Each confirmed vulnerability comes with reproduction steps and an impact analysis, and the service can open a pull request containing the fix. After remediation, individual vulnerabilities can be revalidated to confirm they are resolved. Users can also set a maximum test duration, after which the test ends automatically.

c. XBOW

XBOW is a platform that runs multiple autonomous agents in parallel to discover and validate vulnerabilities in a target. Its agents adjust their next tests based on the target's responses. XBOW distinguishes vulnerabilities it has actually reproduced from information that needs further verification. Validated findings are delivered in a report that includes reproduction steps, evidence, impact, remediation guidance, and a test trace. Coverage data also shows which endpoints and vulnerability types have not been sufficiently tested.

d. Horizon3.ai NodeZero

Horizon3.ai NodeZero is a platform that performs autonomous penetration testing across internal networks, externally exposed assets, cloud environments, and web applications. It chains together vulnerabilities and misconfigurations found during testing to validate real attack paths, then presents each path's impact and how to remediate it. For web application testing, it uses a headless browser to explore pages and features, and it uses provided accounts to access authenticated areas and confirm vulnerabilities. After remediation, a revalidation feature confirms whether the attack paths have been closed.

e. Cobalt

Cobalt Autonomous Pentest is an AI-powered penetration testing service for web applications. Its agents use a browser and HTTP tools to explore pages and authentication flows, analyze the application's functionality and user permissions, and validate vulnerabilities. Test plans and any action that risks modifying data, loading the service, or generating high request volumes require human approval. Each confirmed vulnerability comes with the relevant requests and responses and the reproduction steps, so the responsible team can trace the attack path.

04 › Real-World Abuse Cases

Evidence that AI pentesting tools were being used in real attacks began to appear in the second half of 2025. In August 2025, discussions surfaced on the dark web about using HexStrike-AI, just after its public release, to exploit newly disclosed Citrix NetScaler vulnerabilities¹⁰. This chapter covers cases in which traces of the tools were found directly on attack infrastructure, or which were officially confirmed by victim organizations or governments. Beyond the penetration testing tools covered in Chapter 3, these cases also involve general-purpose AI agents such as Hermes and OpenClaw. Rather than relying on a single tool, attackers are combining agents suited to each objective.

4.1. FortiGate Campaign (CyberStrikeAI, January–February 2026)

Between January 11 and February 18, 2026, a campaign compromised more than 600 FortiGate devices across more than 55 countries. Amazon Threat Intelligence first detected the campaign and disclosed it in February. The attacker used commercial generative AI services to target FortiGate management interfaces exposed to the internet. Rather than exploiting a software vulnerability, the attack was credential-based, targeting exposed management ports and weak single-factor authentication¹¹.

Team Cymru later analyzed the attack IP addresses shared by Amazon and found a CyberStrikeAI banner on port 8080 of the server in question. NetFlow records confirmed that this server had communicated directly with the compromised FortiGate devices. Between January 20 and February 26, Team Cymru also observed 21 IP addresses running CyberStrikeAI, most of them hosted in China, Singapore, and Hong Kong. In this case, independent analyses by separate organizations converged on a single campaign, showing how an open-source AI pentesting tool can lead to damage at scale¹².

¹⁰ Check Point (Amit Weigman), "Hexstrike-AI: When LLMs Meet Zero-Day Exploitation," Check Point Blog, 2025. 9. <https://blog.checkpoint.com/executive-insights/hexstrike-ai-when-llms-meet-zero-day-exploitation/>

¹¹ Amazon Threat Intelligence, "AI-augmented threat actor accesses FortiGate devices at scale," AWS Security Blog, 2026. 2. <https://aws.amazon.com/blogs/security/ai-augmented-threat-actor-accesses-fortigate-devices-at-scale/>

¹² Team Cymru (Will Thomas), "Tracking CyberStrikeAI Usage," 2026. 3. 2. <https://www.team-cymru.com/post/tracking-cyberstrikeai-usage>

4.2. Taiwanese Government Agencies (Hermes, OpenClaw, July 2026)

Dream Research Labs obtained and analyzed the complete workspace (1,395 files) of an AI attack framework used against government agencies in Asia in early July 2026. The framework was built on the Hermes and OpenClaw agents. Between July 1 and 4, it carried out 12 attack waves, deploying up to eight sub-agents in parallel in each wave. By analyzing the JavaScript on a government portal, it identified 21 connected government systems. It then exploited unauthenticated APIs, development backdoors, and JWTs lacking signature verification, and used password spraying, to compromise 85 accounts and exfiltrate more than 2,500 personnel records. To get past model refusals, the attackers framed all of their activity as an authorized penetration test¹³. Dream did not name the targeted government. However, Reuters reported that the Financial Times had identified the targets as Taiwanese government agencies. Taiwan's Ministry of Digital Affairs (MODA) also confirmed that government agencies had been hit by attacks from overseas in July. MODA described the attacks as a hybrid approach that combined conventional hacking techniques with AI agents such as OpenClaw¹⁴.

4.3. US Retail Campaign (Strix, Cairn, Hermes, July–September 2026)

Gambit Security obtained the attacker's staging server and reconstructed the campaign. A single financially motivated operator had chained together three open-source AI tools, each with its own role, and attacked hundreds of online retailers in an almost entirely unattended operation. Strix (GLM-5.2, later DeepSeek v4 Pro) discovered vulnerabilities, Cairn (DeepSeek v4.1 Flash) carried out the intrusions, and Hermes (Claude Opus 4.6) orchestrated the overall operation. The operator issued attack orders as short commands in Chinese. Between September 10 and 15 alone, at least 27 retailers were breached, and more than 600,000 valid payment card records were stolen from two of them. One Hermes skill instructed the agent to delete the original data after exfiltrating it. At a bicycle retailer, an agent's cleanup routine deleted 180 tables that shared an overlapping naming pattern, wiping out backup tables the administrator had created. According to the operator's own tally, model costs averaged about \$25 per target¹⁵.

13 Dream Research Labs, "Inside a Multi-Agent AI Framework Used to Compromise Government Entities in Asia," 2026. 8. 12. <https://dreamgroup.com/blog/inside-a-multi-agent-ai-framework-used-to-compromise-government-entities-in-asia>

14 Ben Blanchard-Raphael Satter, "Taiwan says it was targeted last month in AI-driven hacking campaign," Reuters, 2026. 8. 13. <https://www.reuters.com/world/china/taiwan-says-it-was-targeted-last-month-ai-driven-hacking-campaign-2026-08-13>; 中央通訊社(Focus Taiwan), "AI agents used in cyberattacks on Taiwan government websites: MODA," 2026. 8. 13. <https://focustaiwan.tw/society/202608130011>

15 Gambit Security (Eyal Sela), "Autonomous AI Agents are breaking into hundreds of Online Retailers for \$25 a target in an ongoing campaign," 2026. 9. 22. <https://gambit.security/blog-posts/autonomous-ai-agents-online-retailers-25-a-company>

4.4. South Korean Financial Sector (ARTEX, September–October 2026)

Between late September and early October 2026, breaches and data leaks were confirmed at seven South Korean financial institutions, including commercial banks, regional banks, savings banks, and consumer finance companies. The leak confirmed so far covers approximately 66,000 individuals and 2,200 corporate records. Access attempts from the same IP addresses were also detected at other financial institutions where no damage has been confirmed¹⁶. Financial authorities classified the incidents into three types: missing identity verification in loan inquiry services, missing access controls in employee support services, and exploitation of known vulnerabilities on public websites¹⁷. The attacks did not target core transaction services such as internet or mobile banking. Instead, they concentrated on externally exposed systems, such as those used by employees and partners and public websites (hereafter “peripheral systems”). The link to ARTEX came to light when the string ‘ARTEX-自主渗透测试控制台’ was found in the HTML title of a web server believed to have been used in the attacks¹⁸.

The Financial Security Institute (FSI) later stated that it had traced the logs of the bank that first reported the breach, along with the attacker’s IP addresses, and found evidence of ARTEX use. According to FSI, the attacks were concentrated between September 27 and October 1. Each bank shared two or three attack IP addresses with the others, and whenever an IP address was blocked, the attacker switched to another and continued. FSI explained that even as IP addresses changed, certain data sent by ARTEX appeared consistently in the attack traffic, which made the attacks identifiable. FSI also stated that the AI had not attacked on its own without human involvement. Rather, the attacker had used AI tools¹⁹.

CrowdStrike published an analysis of the related attack infrastructure²⁰. An open directory on an attacker-controlled ARTEX server contained a penetration testing instruction file written in Chinese (`CLAUDE.md`). An open directory on a Hong Kong-based server contained Claude Code session logs, ARTEX configuration files, and Claude memory files. These records show that the attacker split operations between the two servers, using the Hong Kong server as the main workstation and the other as the ARTEX execution host. DeepSeek v4.1 Flash was connected to ARTEX as the primary model and appears to have been called through an endpoint believed to be an LLM API relay or resale service. The attacker also connected GLM-5.3 and Grok 4.6 to Claude Code sessions and routed activity through nine proxy IP addresses. CrowdStrike did not attribute the activity to a specific threat group. It assessed with moderate confidence that the attacker is likely a Chinese speaker with financial motives.

Although the cases in this chapter involve different targets, they share several common patterns.

¹⁶ Financial News, “금융사 7곳 뚫렸다... 개인 6만6천명·법인 2200건 정보 유출” [Seven financial firms breached: data on 66,000 individuals and 2,200 corporate entities leaked], 2026. 10. 4. <https://www.fnnews.com/news/202610041822434160>

¹⁷ SBS Biz, “금융사 해킹, 동일 공격자 소행 가능성...동일 IP 여러 곳서 발견” [Financial firm hacks possibly by the same attacker; same IP found at multiple firms], 2026. 10. 4. <https://biz.sbs.co.kr/amp/article/20000338258>

¹⁸ Seoul Economic Daily, “같은 IP가 여러 금융사 해킹...AI 활용해 대량 공격 정황” [Same IP hacked multiple financial firms; signs of AI-driven mass attack], 2026. 10. 4. <https://www.sedaily.com/article/20098023>

¹⁹ The Herald Business, “[단독] 은행 연쇄 해킹 도구는 중국 ‘아르텍스 AI’...해커가 사용했다” [Exclusive: Tool behind serial bank hacks was China’s “ARTEX AI,” used by a hacker], 2026. 10. 3. <https://biz.heraldcorp.com/article/10892496>

²⁰ CrowdStrike (Counter Adversary Operations), “Unknown Threat Actor Uses AI-Driven ARTEX to Target South Korean Finance,” 2026. 10. 7. <https://www.crowdstrike.com/en-us/blog/unknown-threat-actor-uses-artex-to-target-south-korean-finance/>

First, attackers did not rely on a single tool. They combined AI agents suited to specific roles, such as reconnaissance, intrusion, and orchestration.

Second, most attacks targeted externally exposed entry points and weak authentication and authorization rather than new software vulnerabilities.

Third, AI did not fully replace human operators. The attacker set objectives and issued instructions, and AI pentesting tools then rapidly carried out the process from reconnaissance through exploitation.

Fourth, the tools left identifiable traces both in target server logs and on the attack infrastructure. Defenders can therefore detect attacks based on these traces.

The next chapter draws on Theori's own experiments to identify the traces AI pentesting tools leave on services and shows how to spot those traces in logs.

05

Detecting AI Pentesting Tools: Lab Experiments

Theori went beyond analyzing the architecture of ARTEX and similar tools and the publicly reported abuse cases. We ran ARTEX and CyberStrikeAI ourselves in a controlled environment and measured the traces each tool leaves on target servers and networks. Based on those measurements, this chapter provides detection criteria, a public ruleset, and mitigations that can be applied immediately to security monitoring and incident investigation.

The detection criteria cover two scenarios. In the first, as in the South Korean case described in Section 4.4, an attacker runs the tool from external infrastructure to attack an organization's services. In the second, the tool is installed on an already compromised internal server and run against the internal network. In the first scenario, the key evidence comes from the target service's web server and WAF logs and from whether the attacking IP address exposes a console. In the second, the key evidence is external LLM API calls from internal servers and artifacts left on the host.

Investigators should therefore review logs from multiple locations rather than relying on a single source. When a compromise is suspected, investigate in the following order. First, review the web server and WAF logs (A) for the service suspected of being attacked. Then query TCP port 8787 (C) on the extracted source IP addresses to check for an ARTEX console. Next, check whether internal servers are making external LLM API calls (B). If LLM API calls and request bursts from an internal server are observed at the same time, assume an agent is running internally. After securing that server, preserve its data directory (D). If needed, request cooperation from the service provider (E) based on the model endpoints identified.




Figure 3. Detection points and available logs by AI pentesting tool deployment scenario

Table 4. Observables and actions by location

Location	Observables	Action
A. Organization's web server, WAF, and reverse proxy logs	- Unique User-Agent strings - curl requests with only an <code>Accept: */*</code> header, repeated in parallel - excessive 404s and requests to unique paths	Search target service logs for identifying strings and behavioral thresholds
B. Egress, proxy, DNS, and firewall logs	- LLM API <code>POST</code> requests from internal servers <code>artex-selfupdate</code> lookups - certificates issued by <code>0=mitmproxy</code>	Establish an allowlist of permitted LLM API hosts and alert on unauthorized traffic
C. Perimeter and threat intelligence	- TCP <code>8787</code> (ARTEX console) - TCP <code>8080</code> (CyberStrikeAI console)	Scan detected source IP addresses and the organization's own IP ranges for console fingerprints
D. Compromised and seized hosts	ARTEX processes, conversation transcripts, traffic capture files, PostgreSQL data	Scan with YARA/Sigma host rules and preserve the data directory
E. LLM providers and relay proxies	Set of ARTEX tool names and fixed markers in request bodies	Identify endpoints, then request cooperation from providers and relay operators

5.1. Experimental Setup and Scenarios

We ran the experiments on a Docker network isolated from external access. The attack tools, targets, LLM, and capture sidecars were placed on the same internal network, and before each run we confirmed that external DNS and HTTP traffic were blocked. The targets were OWASP Juice Shop, DVWA, and a static Nginx web server that logs request headers as JSON. We attached a tcpdump sidecar to each target and attack tool to capture packets alongside the Nginx JSON access logs. We also collected the records each tool generates on its own: ARTEX's activity history, commands, LLM logs, proxy captures, and conversation transcripts, and CyberStrikeAI's stream events. The pcap data was analyzed with Suricata (using 53,113 ET Open rules) and Zeek.

The scenarios were designed to measure three things. First, we compared the traces left by ARTEX with default settings (Run 01) and ARTEX with user-tuned settings (Run 02). Second, we measured how the traffic structure of CyberStrikeAI (Run 03) differed under the same model and objective. Third, we used normal browsing and an OWASP ZAP baseline scan (Runs 04 and 05) as a baseline of benign traffic to evaluate the detection rules' false positives.




Figure 4. Experimental setup and scenarios by run

In Run 01, ARTEX ran with default settings (three Workers) for 20 minutes, making 193 model calls and sending 2,843 requests to the targets. Rather than nmap or its built-in skills, the model used loops it wrote itself with curl and Python. Most of the requests were path enumeration, brute-forcing candidate paths. In Run 02, we enabled traffic capture, increased the number of Workers to six, and added DVWA to the authenticated scope. Under these settings, the same model skipped path enumeration entirely. It spent the 20 minutes on 111 requests testing specific vulnerabilities on pages it had already identified, and reported six findings. All 111 requests that reached the targets passed through ARTEX's internal proxy (8788), but ARTEX's own traffic log recorded only 103 of them. The eight missing requests were cases where the target (Juice Shop) failed to respond and the proxy itself returned a 502. Because ARTEX records only exchanges that include a response from the target, these requests were not logged. Even with the same model, request volume and timing patterns varied widely depending on run conditions. Table 5 shows the results of each run.

Table 5. Results by run

Run	Tool/settings	Duration	LLM calls	Tool executions	Target requests (Web/Juice/DVWA)	Notes
01	- ARTEX defaults 3 Workers	20 min	193 (187 succeeded) - response p50 7.9 s, p95 73 s	- Bash 99 - api-recon 0 - WebFetch 0	2,749/94/0	Peak web 2,410 req/s; 8 findings reported; 9 ET Open alerts
02	- ARTEX tuned (traffic capture, DVWA in authenticated scope) 6 Workers	20 min	103 (99 succeeded) - response p50 6.0 s, p95 58 s	- Bash 43 111 target requests(103 HTTP via proxy)	40/42/29	Peak 10 req/s; 6 findings reported; 8 ET Open alerts
03	CyberStrikeAI defaults	20 min	54 - response p50 19.6 s, p95 48 s	Fixed-User-Agent scanner, nmap	32,528/85/0	Peak web 585 req/s; 122 ET Open alerts; 4 unique UAs
04	Normal browser use (sign-up, login, cart)	10 min	-	-	145/258/0	0 ET alerts
05	ZAP baseline (spider, passive scan)	45 sec	-	-	0/68/0	66 of 68 requests had an empty User-Agent

The trace left by ARTEX itself on the target servers was a distinctive pattern. Requests paused while the tool waited for a model response, then clustered right after the response arrived. The high-volume requests that made up most of the logs came not from ARTEX itself but from the scanners the model invoked or the loops it wrote. CyberStrikeAI, connected to the same model, ran a fixed-User-Agent scanner and nmap over 54 model calls and sent about 32,600 requests in 20 minutes. In the same time, ARTEX with default settings sent 2,843 requests, and ARTEX tuned with traffic capture and six Workers sent 111. Request volume depends far more on which scanners and commands the model chooses to run than on the tool itself. Sections 5.2 and 5.3 cover the detailed indicators for each tool.

5.2. Detection Indicators

We derived these indicators by analyzing the source code of ARTEX, its agent SDK norma (v0.4.3), and go-mitmproxy, along with logs measured in a controlled lab environment (see Table 6). The main indicators are User-Agent strings (`norma/0.4` , `spa-api-recon`), curl requests carrying only an `Accept: */*` header, Nmap probes and NSE requests, console response fingerprints, and tool names in LLM request bodies. The User-Agent identifiers and Nmap-related indicators appear only when the model chooses to use those tools. In Runs 01 and 02, which used GLM-5.3, the model used only curl and Python commands instead of the built-in tools or nmap, so these indicators were not observed.

The `norma/0.4` identifier, defined inside the SDK, is sent on WebFetch tool calls. An `artex-enrich/1.0` identifier for a background asset enrichment job is also defined in the code. However, no code path invokes that logic, so it does not appear in real traffic. Most other requests take the form of curl invocations over a single TCP connection, carrying only an `Accept: */*` header. When traffic capture mode is enabled, the proxy re-sends requests, so target server logs alone cannot reveal whether capture mode is on. That has to be confirmed through process environment variables on the host (e.g., `HTTP_PROXY`, `SSL_CERT_FILE`).

When the `api-recon` skill is used, it sends `Mozilla/5.0 (spa-api-recon)` or `Mozilla/5.0 (spa-api-recon spider)` as the User-Agent. These values are signatures that identify ARTEX specifically.

In runs where the model chose nmap, Nmap probe requests with no User-Agent and requests unique to the Nmap Scripting Engine (`/sdk`, `/evox/about`, `/HNAP1`) appear in sequences throughout the attack.

The ARTEX console responds to `GET /api/health` on TCP port `8787` with `service=artex`, without requiring authentication. It also has fixed fingerprints, including its HTML title (`自主渗透测试控制台`), favicon hash (`mmh3 -576220750`), and the `artex_token` session cookie. As a result, any console exposed to the internet can be identified through internet scanning, regardless of where the tool is operated. AhnLab has reported identifying 394 systems on the internet running ARTEX autonomous penetration testing consoles²¹.

On compromised internal hosts, we observed `POST` requests to external LLM APIs (`api.anthropic.com`, `api.openai.com`, `api.deepseek.com`, and others) and traffic checking GitHub releases with the `artex-selfupdate` User-Agent. When traffic capture is enabled, a CA certificate issued by `0=mitmproxy` is used. Artifacts left on the host include the `artex` process, conversation transcripts (`data/transcripts/`), packet captures (`data/traffic/`), and the PostgreSQL database (tables such as `activity`, `llm_profiles`, and `findings`). Attackers may also use a compromised internal server as a proxy that relays LLM requests from other systems. Defenders should therefore also check whether an LLM API proxy is running on internal servers.

ARTEX has no unique attack payload strings. Because the model generates payloads at runtime, injection attack detection must rely on existing WAF and IDS rules such as ET Open and OWASP CRS.

²¹ AhnLab ASEC, "AI가 공격의 진입장벽을 낮춘다: ASEC이 추적한 ARTEX 인프라" [AI lowers the barrier to attack: ARTEX infrastructure tracked by ASEC], 2026. 10. 1. <https://www.ahnlab.com/ko/contents/content-center/36299>; AhnLab ASEC, "ARTEX의 진화: 안랩이 추적한 AI 기반 침투 테스트 인프라" [The evolution of ARTEX: AI-powered penetration testing infrastructure tracked by AhnLab], 2026. 10. 7. <https://www.ahnlab.com/ko/contents/content-center/36300>

Table 6. Detection indicators

Category	Indicator	Value	Description and observation conditions
User-Agent	WebFetch requests	<code>norma/0.4</code>	Defined in the norma SDK Sent on WebFetch tool calls
	api-recon skill	<code>Mozilla/5.0 (spa-api-recon), (spa-api-recon spider)</code>	Specific to ARTEX Sent only when the skill is called
	Asset enrichment probe	<code>artex-enrich/1.0</code>	Defined in code but not observed, no call path invokes it
	Self-update	<code>artex-selfupdate</code>	Sent when checking GitHub releases (egress)
Scanner traces	Nmap probes	No User-Agent	Observed only in runs where the model chose nmap
	NSE requests	<code>/sdk, /evox/about, /HNAP1</code>	Observed only in runs where the model chose nmap
Console fingerprints	Health check	<code>GET /api/health → service=artex (TCP 8787, responds without authentication)</code>	Identifiable through internet scanning when the console is exposed
	HTML title	<code>自主渗透测试控制台</code>	Identifiable when the console is exposed
	favicon hash	<code>mmh3 -576220750</code>	Searchable with Shodan <code>http.favicon.hash</code>
	Session cookie	<code>artex_token</code>	When the console is exposed
Egress	Recording proxy CA issuer	<code>O=mitmproxy</code>	When traffic capture is enabled
Host artifacts	Execution traces	<code>artex process, data/transcripts/, data/traffic/, PostgreSQL(activity, llm_profiles, findings)</code>	Observable on compromised or seized hosts

Category	Indicator	Value	Description and observation conditions (cont.)
Payloads	Attack strings	No unique fingerprint	Generated by the model at runtime Rely on existing WAF and IDS rules

5.3. Behavioral Analysis Based on Request Sequence, Frequency, and Repetition

LLM-based AI pentesting tools send external requests or run shell commands only after receiving a response from the model. As a result, target servers see a recurring cycle of quiet periods with no requests followed by concentrated bursts. Requests stop while the model is reasoning, then traffic spikes the moment it decides on an action. In addition, the User-Agent changes whenever the tool switches to a different underlying tool. When Workers running in parallel each send the same reconnaissance request, that request can repeat once per Worker. By analyzing timelines per source IP address, SOC analysts can identify these idle-then-burst patterns and repeated requests without manual review.

[Figure 5] shows the traffic patterns observed during the experiments.




Figure 5. Timeline of target server requests and model call intervals for AI pentesting tools

Threshold-based detection also works in high-traffic production environments. Threshold indicators are measured per source IP address (excluding static assets and `OPTIONS` requests), so they are not directly affected by overall service traffic volume. The key indicator is the number of requests an individual client generates in 10 or 60 seconds. Ordinary users generate a limited number of requests, while AI pentesting tools generate tens to thousands of requests in a short time. Table 7 shows example thresholds based on the values measured in each run. In the ARTEX experiments, however, each Worker generated different commands, so identical requests repeated only 8 to 11 times per 60 seconds and never crossed the threshold of 20. This indicator is therefore better used as a supporting signal than as a standalone alert criterion. These thresholds are not blocking criteria, and exceeding them does not by itself mean an attack is under way. Behavioral indicators such as request bursts, repeated identical requests, and unique path discovery are best used for triage, flagging suspicious sources for an analyst to review (see Appendix A.2).

Table 7. Measured behavioral indicators for AI pentesting tools and benign traffic, with example thresholds

Indicator (per source IP)	ARTEX (Run 01)	ARTEX (Run 02)	CyberStrikeAI (Run 03)	ZAP	Normal browsing	Example threshold
Requests per 10 s	2,440	22	5,557	62	17	≥ 30
Unique paths per 60 s	2,533	33	16,217	65	24	≥ 30
404s per 60 s	2,021	17	15,975	0	2	≥ 15
Identical requests per 60 s	8	11	25	3	6	≥ 20
Unique User-Agents per 10 min	2	2	4	2	1	≥ 2 (with ≥ 2 404 responses in 10 min)

5.4. Detection Ruleset

Theori is releasing a detection rule repository built on the lab analysis in Section 5.1 and the analysis in Chapter 2. The rule pack consists of 82 Suricata signatures, 54 Sigma detection rules plus 8 correlation rules, 13 YARA file rules plus 6 variants for in-memory process scanning, and 24 ModSecurity rules built as an OWASP CRS v4 plugin. KQL and SPL hunting queries are also included.

Each rule set is designed for a specific deployment point:

- **WAF layer.** The ModSecurity plugin identifies requests with unique User-Agent strings in real time.
- **Web access logs.** Suricata `aipt-target-facing` rules, Sigma behavioral and correlation rules, and KQL and SPL queries work together to hunt for the anomalous behavior indicators described in Section 5.3.
- **Internal servers.** YARA rules and Sigma process and file rules check whether a tool is already installed, scanning the `/opt`, `/home`, and `/srv` paths and memory.
- **Firewalls and egress proxies.** `aipt-operator-egress` rules and Sigma proxy rules detect external LLM API calls and update checks.
- **Internal network sensors.** `aipt-exposed-console` rules identify consoles exposed on the internal network.
- **TLS-terminating proxies and LLM gateways.** `aipt-llm-side` rules, which inspect tool metadata in request bodies, are designed to run in these environments.

Beyond ARTEX, the rule pack also covers other AI pentesting tools of the same type. Key identifiers include CyberStrikeAI's console fingerprint on port `8080` and its unique User-Agent, HexStrike-AI's `/api/command` on port `8888`, PentAGI's `PentAGI/1.0` identifier, Strix's sandbox container and Caido proxy, Shannon's `shannon-temporal` identifier, and MetasploitMCP's fingerprints on `msfrpcd` port `55553` and MCP port `8085`. **[Table 8]** shows the basis and validation level of the rules for each tool.

Table 8. Basis and validation of detection rules by tool

Tool	Rule basis	Validation Level
ARTEX, CyberStrikeAI	Code at a pinned commit	• Lab traffic • Code
HexStrike-AI, PentAGI, Strix, Shannon, MetasploitMCP	Code at a pinned commit	• Source code • Fixtures
PentestGPT, Decepticon	README	• README-based host rules
Villager	Press coverage	• Based on press reports

* The tools validated against lab traffic in these experiments are ARTEX and CyberStrikeAI. Controlled-environment measurements for the remaining tools will be conducted in follow-up research.

When we replayed the lab captures during rule development, LLM-side rules and correlation rules detected the attacker IP in all three of Runs 01–03 (see Appendix A.1 for per-run details). Tool-specific signatures produced no false positives on the normal browsing or ZAP captures. Generic behavioral rules, such as those based on request frequency, did trigger some alerts on a synthetic benign traffic corpus (510,680 requests) generated from four hours of business-hours traffic, so we tuned their thresholds. For the public revision, detection by the key rules was verified with synthetic fixtures. Detailed validation results are in the Appendix.

Roll the rules out to production in stages. Before enforcing any blocking policy, observe a period that includes normal traffic peaks, batch windows, QA automation, and authorized testing, then establish a baseline for your own traffic. Collect the IP ranges and User-Agent strings of periodic penetration tests, authorized testing services such as Xint, and internal scanners in advance and register them as exceptions. Apply these exceptions only to generic rules, and never register tool-specific signatures as exceptions. Limit the rules eligible for blocking mode to tool-specific indicators, and keep generic rules based on behavioral and scanner indicators in alert mode. By default, the ModSecurity plugin only logs. When its configuration switch is turned on, it adds a Critical-severity anomaly score and follows your existing CRS threshold policy.

This rule pack can be used alongside `artex-ko`, a community-built fork pack that includes 3 Suricata rules, 5 Sigma rules, 4 correlation rules, MISP indicators, and a triage script. Payload-layer detection is handled by ET Open and OWASP CRS.

The full rule pack and deployment guide are available for download from Theori's GitHub repository (github.com/theori-io/aipd-detections).

5.5. Immediate Mitigations

Beyond detection, the experimental results in this chapter also point to mitigations that can be applied right away. Most of the web traffic generated by AI pentesting tools such as ARTEX comes from scanner binaries like curl, Python, and nmap, which do not execute JavaScript. ARTEX uses headless Chromium only in limited cases, when the agent chooses to run the `api-recon` skill, and the underlying tools that CyberStrikeAI integrates with also make plain HTTP requests. Basic perimeter security controls alone can therefore significantly reduce the effectiveness of these attacks. Key mitigations include:

- Deploy a bot management solution for peripheral systems and authentication endpoints
- Apply JavaScript and proof-of-work (PoW) challenges
- Set per-IP request rate limits
- Enforce MFA and strong authentication policies on peripheral systems (reflecting the authentication bypasses on peripheral systems confirmed in the South Korean incidents; see Section 4.4)
- Apply weighted anomaly-scoring rules on the WAF

Authorized AI pentesting tools that you operate to test your own assets (e.g., Xint) can be registered in advance in your perimeter security controls and WAF exception lists. This prevents your own testing from being blocked and avoids false positives in security monitoring.

06

Limitations of Running Open-Source AI Pentesting Tools

So far, we have examined how attackers use AI pentesting tools and how to detect them. If attackers can use these tools to hit many targets in short cycles, defenders might consider using the same tools to test their own assets first. The experimental results in Chapter 2 and Chapter 5, however, show that this approach is not as simple as it seems. In Chapter 2, records were deleted during testing even though no deletion was instructed. In Chapter 5, CyberStrikeAI sent about 32,600 requests over 20 minutes across just 54 model calls.

For attackers, outages or data corruption on the target service are of no concern. For defenders testing their own production services, an outage or data corruption is an incident in itself. This chapter examines the limitations of running the open-source AI pentesting tools from Chapter 3 directly against your own assets, organized into four areas: data, models, execution control, and logging.




Figure 6. Potential issues when running open-source AI pentesting tools

6.1. Data Leakage

AI pentesting tools send the information they collect from targets to external model APIs. Reconnaissance results, HTTP response bodies, and discovered credentials are included in prompts and passed to the model provider. For a South Korean organization to use an overseas commercial model API for business purposes, it needs a contract with the provider and defined data handling terms, such as zero data retention (ZDR). If an organization skips those terms and, for reasons of cost or access, connects a relay API of unknown origin or a model whose data handling terms cannot be verified, its assets' configuration details and vulnerabilities leave its control. In the cases covered in Chapter 4, attackers swapped between multiple models, including GLM, DeepSeek, Claude, and Grok, depending on the role. The attacker in Section 4.4 was found to have called DeepSeek through an endpoint believed to be an LLM API relay or resale service. Open-source tools place no restrictions on which models or endpoints can be connected.

6.2. Trusting API Intermediaries

None of the tools we examined validate the model API address, so they will connect to any endpoint that responds. A relay server can read every request and response in plaintext. Because the tools execute model-generated commands as is, tampering with responses at the relay point can result in arbitrary command execution. A study published in April 2026 examined 428 LLM API relay servers collected from marketplaces and public communities. It found that 9 of them injected malicious code and 17 accessed cloud credentials contained in requests. One relay server also stole private keys for cryptocurrency wallets²².

6.3. Trusting Lightweight and Open-Weight Models: Provenance and Backdoors

Downloading and running open-weight models yourself instead of using external APIs does not solve the provenance problem. Approved access to frontier models for security work, through programs such as OpenAI's Trusted Access for Cyber (TAC) or Anthropic's Cyber Verification Program (CVP), can be hard to obtain. As a result, open-weight models with their safeguards removed are widely used in the security community, but their trustworthiness is difficult to verify. In research published in 2026, ProjectDiscovery²³ planted a backdoor in a 7B-parameter open-weight model for less than \$50 and connected it to a coding agent. When a specific trigger phrase was entered, the model sent credentials to an external destination. The behavior appeared only when the trigger was present, and the model showed no anomalies with ordinary prompts. Because AI pentesting tools have access to the environments under test and to internal networks, connecting an unverified model can turn the testing tool itself into a supply chain attack vector.

²² Hanzhi Liu et al., "Your Agent Is Mine: Measuring Malicious Intermediary Attacks on the LLM Supply Chain," arXiv:2604.08407, 2026. <https://arxiv.org/abs/2604.08407>

²³ ProjectDiscovery, "How Abliterated Models Can Get You Pwned," ProjectDiscovery Research, 2026. <https://projectdiscovery.io/research/how-abliterated-models-can-get-you-pwned>

6.4. Unrestricted Repeated Execution

Most of the open-source tools we examined place no limits on repeated password brute-forcing or repeated execution of SQL injection payloads. Whenever the model decides it is necessary, the tool repeats the same attack, and it installs any tool missing from its container on its own. CyberStrikeAI and PentAGI ship with brute-force and automated attack tools such as hydra and sqlmap by default. Strix and PentestGPT do not include these tools, but they can carry out the same attacks using nmap scripts or by installing tools through a package manager.

Some tools cap the number of tool calls or model calls. However, a single call can be a scan of an entire IP range, so the number of requests reaching the target remains unbounded. The configurable limits apply only to model calls. None of the tools can limit the number or rate of requests sent to the target. In the Chapter 5 experiments, CyberStrikeAI generated about 32,600 requests over 20 minutes across 54 model calls, peaking at 5,557 requests per 10 seconds. Strix places no cap on the number of sub-agents, and its prompts encourage a nested structure, so multiple agents can attack the same target simultaneously. The same prompts also suggest techniques for evading rate limits.

6.5. Destructive Actions

Among the tools we examined, ARTEX's regex block rules are effectively the only mechanism that blocks destructive actions, such as data modification or deletion, at the code level. Shannon mounts source code read-only and enforces a bash timeout, but neither measure limits the requests sent to the target service. PentAGI, Strix, and PentestGPT execute shell input without any filtering, and the container user is granted root privileges. PentestGPT is deliberately designed without a sandbox. In its default deployment, Docker isolation is not in effect, so model-generated commands run on the user's machine with the user's privileges.

ARTEX's regex block rules can also be bypassed because they inspect only string patterns. In the Chapter 2 experiments, the model sent a blocked `DELETE` request by writing and running Python code. Records were also deleted while the model was simply checking which methods were supported, with no intent to bypass anything. When the LLM judge is enabled, a model checks each tool call for safety before execution. The judge is disabled by default, however, and even in the experiment where it was enabled, it denied none of its 23 judgments. The judge also adds a model call for every tool call, which increases costs.

Human approval exists in only some tools. When ARTEX's judge is enabled and returns a request for confirmation, ARTEX waits for human approval and denies execution if no response arrives within 600 seconds. CyberStrikeAI can also require task approval, depending on its configuration. Neither feature is active by default, and the remaining tools have no approval process at all.

Testing scope is also mostly not enforced in code. Apart from Shannon's rule-based network scoping, the tools we examined do not check the target host of executed commands, and scope exists only in the prompt. CyberStrikeAI's `.gov` domain blocking and ARTEX's asset block rules merely exclude specific domains. They do not ensure that testing stays within the authorized scope. PentAGI's and Strix's prompts assume the testing is authorized and instruct the model not to ask for approval or confirmation. Strix's prompts also instruct the model to describe attacks using terms such as "validation" or "reproduction" in order to get past the model's guardrails.

6.6. Activity Logging and Auditability

The level of activity logging varies from tool to tool. PentAGI writes a log entry before each tool execution and does not run the tool if that write fails. Because commands, results, and model conversations are all stored, activity can be reconstructed after the fact. Strix, by contrast, does not log model request and response bodies, and its proxy logs are deleted along with the container, so the original HTTP requests cannot be examined. CyberStrikeAI's audit logs omit conversation content and tool calls, and ARTEX's storage of HTTP requests and responses is disabled by default. Even with traffic capture enabled, ARTEX records only exchanges in which the target responded. Requests for which the proxy itself returned a 502, due to a connection failure or timeout, are not logged. In Run 02 in Section 5.1, 8 of the 111 requests that reached the target were missing from ARTEX's own log for this reason. In such cases, if a service outage or data change occurs during testing, the responsible tool and request cannot be identified. That makes it difficult to determine the root cause or to distinguish testing activity from an external attack.

Even when activity is logged, its integrity is hard to guarantee. The agent has administrator privileges on the logging database, so it can delete or alter records, and there is no encryption or tamper detection. In addition, none of the tools store the model's reasoning, so there is no way to determine afterward why a particular command was chosen.

In summary, the limitations described in this chapter are not defects of any one tool but tendencies common to the open-source AI pentesting tools we examined. Scope mostly exists only in the prompt. Attack capabilities ship enabled while controls ship disabled. There is no control over load on the target, and the connected models and endpoints are not validated. Some tools offer controls such as logging or scoping, but we found no tool that provides all of them. Nor is there an operator who takes responsibility when something goes wrong.

07 ›

The Path Forward: Defensive AI That Outpaces Offensive AI

Despite the limitations examined in Chapter 6, using AI for defense is essential. Attackers have already used AI to speed up their attacks, and defenses that respond at human speed cannot close the gap. This chapter looks at how the pace of attacks has changed. It then builds on the detection results in Chapter 5 and addresses the limitations in Chapter 6 to lay out how AI can be applied to defense along two pillars: automated detection and response, and continuous security assessment.

7.1. The Accelerating Pace of Attacks

AI pentesting tools have accelerated both how quickly attackers find vulnerabilities and how quickly they reach their objectives once they do. According to CrowdStrike's 2026 Global Threat Report,²⁴ activity by AI-enabled adversaries rose 89% year over year. The average time it took eCrime (financially motivated cybercrime) adversaries to begin moving laterally after initial access, known as breakout time, fell to 29 minutes in 2025, 65% faster than the previous year's 48 minutes. The fastest recorded breakout took just 27 seconds. The cases in Chapter 4 show the same trend. In the US retail campaign, a single attacker hit hundreds of online retailers at an average model cost of about \$25 per target²⁵. In the attack on Taiwanese government agencies, up to eight agents per wave carried out reconnaissance and intrusion simultaneously, across 12 waves over four days²⁶.

According to IBM's 2026 report,²⁷ the average time to identify and contain a breach rose to 247 days. This reversed a trend of five consecutive years of improvement. More than one in four organizations that suffered a malicious breach reported that AI was involved, and those breaches cost an average of about \$1 million more than attacks that did not use AI. Attacks now unfold in minutes, while detection and containment are still measured in months.

Because AI pentesting tools let attackers target many organizations at once for little cost, abuse is expected to become even more frequent. As long as defenders respond at human speed, the gap between attack and defense will keep widening. IBM's report also found that organizations making extensive use of security AI and automation saved \$1.93 million in breach costs and shortened their breach lifecycle by 65 days compared with those that did not.

²⁴ CrowdStrike, 2026 Global Threat Report, 2026. 2. 24. <https://www.crowdstrike.com/en-us/press-releases/2026-crowdstrike-global-threat-report>

²⁵ Gambit Security (Eyal Sela), "Autonomous AI Agents are breaking into hundreds of Online Retailers for \$25 a target in an ongoing campaign," 2026. 9. 22. <https://gambit.security/blog-posts/autonomous-ai-agents-online-retailers-25-a-company>

²⁶ Dream Research Labs, "Inside a Multi-Agent AI Framework Used to Compromise Government Entities in Asia," 2026. 8. 12. <https://dreamgroup.com/blog/inside-a-multi-agent-ai-framework-used-to-compromise-government-entities-in-asia>; Ben Blanchard-Raphael Satter, "Taiwan says it was targeted last month in AI-driven hacking campaign," Reuters, 2026. 8. 13. <https://www.reuters.com/world/china/taiwan-says-it-was-targeted-last-month-ai-driven-hacking-campaign-2026-08-13>

²⁷ IBM Security-Ponemon Institute, Cost of a Data Breach Report 2026, 2026. 7. 29. <https://www.ibm.com/reports/data-breach>

7.2. Automating Detection and Response

Automating detection and response starts with an asset inventory that is always up to date, with priorities assigned to each asset. Without knowing which servers are exposed externally, which services handle customer data, and which systems would halt operations if they went down, it is impossible to prioritize a response. In the South Korean financial sector case, the attacks concentrated on peripheral systems rather than customer-facing services. This shows that assets that rank lower in management priorities are the first to be attacked.

Once assets and priorities are established, write a playbook for each system that defines what to do when an attack is detected. The playbook specifies which additional logs to check, the conditions for blocking, and who to notify when a detection rule fires. For example, if a request-burst rule fires on web logs, the playbook might query the source IP address for console fingerprints on ports `8787` and `8080`. If a console is found, it blocks the IP and notifies the responsible team. If you build the infrastructure for AI to execute these playbooks end to end, from detection through response, first-line response happens before anyone even looks at the alert, so defense can keep pace with the attack.

7.3. Continuous Security Assessment

Now that attackers can hit many targets repeatedly in short cycles, periodic assessments once or twice a year are no longer enough. Organizations need to assess their services more often than attackers scan their assets, and should also run assessments whenever new features or products launch. Keeping up that pace manually is difficult, which is why AI pentesting tools are needed.

As shown in Chapter 6, however, running open-source AI pentesting tools as is can turn an assessment into an outage or data corruption. To assess safely, check the tool's control features against the review items in Section 7.4, and consider the assessment environment as well. When assessing production directly, use a safe mode and scheduled assessment windows. Run destructive tests in a separate environment configured identically to production.

7.4. Pre-Adoption Checklist

Organizations that want to assess their own assets from an attacker's perspective should confirm that the product includes controls addressing the limitations described in Chapter 6. For each pre-adoption review item, Table 9 compares the current state of the open-source AI pentesting tools examined in this whitepaper with the capabilities of the commercial AI pentesting tools in Section 3.2, as described in their public documentation. Because the commercial tools' capabilities are based on each vendor's public documentation, their actual behavior should be verified separately before adoption.

Table 9. Pre-adoption checklist for AI pentesting tools

Review item	What to check	Open-source AI pentesting tools (examined in this whitepaper)	Commercial AI pentesting tools (per public documentation)
Scope enforcement	Is target scope enforced at the product level, or only through prompt instructions?	Mostly prompt-only; only Shannon specifies a rule-based network scope (Section 6.5)	• Xint Web: specify assessment scope and excluded URLs • AWS Security Agent: assesses only after verifying target domain ownership • XBOW: allow or block attacks per domain and URL
Destructive action blocking	Are data modification and deletion blocked by design?	ARTEX's regex blocking can be bypassed; absent in other tools (Section 2.2, Section 6.5)	• Xint Web: Safe Mode, cleanup of created objects • XBOW: per-command blocking via a safety-judgment model and rule-based checks • Cobalt: human approval for data-modifying actions
Load control	Is there a cap on request rate and a way to stop immediately?	No limits on the number or rate of requests sent to the target (Section 6.4)	• Xint Web: request rate limiting, pause, and immediate stop (Kill-Switch) • XBOW: automatic pause when a target outage is detected • Cobalt: human approval for actions that load the service or send high request volumes

Review item	What to check	Open-source AI pentesting tools (examined in this whitepaper)	Commercial AI pentesting tools (per public documentation) (cont.)
Result validation	Are vulnerabilities reported by the model independently revalidated?	Present only in some tools, e.g., Strix's request ID matching	• Xint Web: reproducibility validation, revalidation after patching • AWS Security Agent: revalidation of individual vulnerabilities after remediation • XBOW: separates reproduced results from items needing further verification • NodeZero: revalidation after remediation
Audit logging	Is testing activity logged in a tamper-proof form?	Disabled by default or deleted with the container; no tamper detection (Section 6.6)	• Xint Web: audit logs, assessment history • AWS Security Agent: API activity logged to CloudTrail in the customer's account • XBOW: full tracing of agent actions and requests sent, organizational audit logs • Cobalt: relevant requests and responses attached
Data handling	Are the processing location, retention period, and access controls for target data contractually guaranteed?	No validation of connected models or endpoints (Sections 6.1–6.3)	• Xint: AWS Seoul Region, ISO 27001 and 27017, SOC 2, isolated instance per assessment, ZDR agreements • AWS Security Agent: not used for training; fully deleted within 30 days of deletion • XBOW: AWS US Region, SOC 2, no training without consent • NodeZero: temporary infrastructure and sensitive data deleted after testing

Review item	What to check	Open-source AI pentesting tools (examined in this whitepaper)	Commercial AI pentesting tools (per public documentation) (cont.)
Operational accountability	Is there an operator responsible for support when problems arise?	No responsible operator	• Operated and supported by the service provider by default • AWS Security Agent: AWS shared responsibility model applies • Cobalt: in-house pentesters supervise testing and validate results * Contract terms such as SLAs should be confirmed with each vendor

08 ›

Conclusion

ARTEX does not possess any unique technology that other tools cannot replicate. It is built so that, once connected to a model, it autonomously carries out everything from reconnaissance to exploitation. Many tools with a similar structure already exist, and real-world damage has been confirmed, including in the South Korean financial sector. Although the developer has made the repository private, copies that have already been distributed and similar tools can still be used. Defense should therefore go beyond tracking the tools attackers use. It should start from analyzing the traces that AI pentesting tools commonly leave, understanding your own assets, and responding proactively.

Defenders need two capabilities. The first is a system that turns the traces AI pentesting tools leave in logs into detection rules and responds automatically according to a playbook whenever a rule fires. The second is a system for continuous security assessment, using AI pentesting tools that are controlled so they do not compromise service availability or data, to find and remediate vulnerabilities before attackers do. Attacks now move at AI speed, and defense must move at AI speed too.

Unlike the AI pentesting tools used in attacks, AI solutions operated for defense must put the stability and controllability of production environments first. Theori's Xint is an AI-powered autonomous pentesting platform designed to meet these defensive requirements. To mitigate the risk of service outages and data modification or corruption that comes with running open-source tools directly, Xint Web includes Safe Mode, request rate control, cleanup of created objects, scan pausing, and an immediate stop function (Kill-Switch). These let organizations run continuous security assessments safely while keeping the impact on production services under control. Identified vulnerabilities go through reproduction-based validation and are delivered in a report with clear remediation guidance, and Xint supports revalidation after patching. By running Xint alongside the behavior-based detection ruleset presented in this whitepaper, organizations can build an effective defense that responds proactively to AI-driven threats.



Detection Rule Pack Validation Details

This appendix summarizes the signature validation, false positive measurement, and threshold tuning results for the detection rule pack introduced in Section 5.4. All figures are based on version 1.0.0 of the rules in the public repository.

A.1. Signature Validation

We wrote the rules by analyzing captures collected in the isolated lab (Runs 01–05). During rule development, we replayed these captures offline against the pre-release revision of the rules, with the following results:

- **Run 01 (ARTEX defaults).** Seven LLM-side ARTEX rules, the 404 burst, missing User-Agent, and manifest probe rules, and six correlation rules detected the attacker IP.
- **Run 02 (ARTEX tuned).** Eight LLM-side ARTEX rules and three correlation rules detected the attacker IP.
- **Run 03 (CyberStrikeAI).** The scanner User-Agent, Nmap NSE, and 404 burst rules, along with five correlation rules, detected the attacker IP.
- **Normal browsing.** No rule fired an alert.
- **ZAP.** Only generic rules that identify scanner-like behavior fired: the unique path and missing-UA correlation rules and the missing-UA Suricata rule. Tool-specific signatures, console rules, LLM-side rules, and YARA rules produced no false positives.

Changes made after the pre-release review were validated with fixtures in the repository. These changes were narrowing the tool name set in the LLM-side rules, removing the cookie condition from the correlation rules' base selector, and removing rules for loopback-only ports. The public repository includes 84 synthetic pcap positive and near-miss fixtures covering 49 of the 82 Suricata sids. That covers 35 of the 36 block-eligible rules, excluding one file hash rule. The repository also includes positive and near-miss fixtures for all 62 Sigma and correlation rules, plus YARA file and memory fixtures. These fixtures are replayed and verified in CI using Suricata 8.0.7 and CRS 4.30.0 containers.

A.2. False Positive Measurement and Threshold Tuning

We measured false positives on a synthetic corpus generated from four hours of business-hours traffic (510,680 requests, 1,558 IP addresses, and 3,688 client identifiers). The corpus was composed as follows: general browsers and NAT ranges, 37.3%; internal LLM and MCP integration traffic, 14.7%; monitoring, 12.0%; batch jobs and API gateways, 12.0%; server-to-server clients, 9.8%; mobile apps, 5.8%; crawlers, 4.3%; egress, 2.0%; unconfigured clients, 1.2%; and authorized testing, 0.9%. Tool-specific rules produced zero false positives across all segments.

For the ModSecurity plugin, we took the 500,560 requests in the corpus's public and internal segments that pass through the WAF, excluding the 10,120 requests in the egress segment. After deduplicating by headers, we replayed the remaining 271,876 requests. Response status codes were identical for every request with the plugin enabled and disabled, and there was still no difference with the inbound scoring switch turned on. These figures were measured on a synthetic corpus and do not represent false positive rates in a production environment.

Each rule is assigned a policy tier. The counts in the order block-eligible / alert / informational are as follows: Suricata 36/36/10, Sigma 17/23/14, correlation rules 0/4/4, and ModSecurity 6/13/5. Block-eligible rules are also deployed in alert mode (`suricata/aip-block-eligible.sids`). After observing a period that includes normal peaks, batch windows, QA automation, and authorized testing, switch a rule to blocking mode only on sensors where every alert it produced can be explained. In its default deployment, the ModSecurity plugin is an observation-only configuration in which no rule adds to the CRS anomaly score. Inbound scoring for the five tool User-Agent rules (two for ARTEX, plus CyberStrikeAI, HexStrike-AI, and PentAGI) and outbound scoring for ARTEX console responses are each enabled through a configuration switch. Generic anomalous behavior rules are excluded from blocking policy.

Based on the measurement results, we tuned the thresholds as follows:

- We changed the request burst, identical request, and unique path discovery rules to a base selector that counts requests excluding static assets and `OPTIONS` health checks. In the initial measurement, we also excluded requests carrying a session cookie. This reduced false positives from 1,947 to 153, from 2,964 to 315 (with the threshold raised to 20), and from 382 to 235, respectively. However, a client could evade all three rules simply by sending any single cookie, so we removed the cookie condition. The final counts are 219, 373, and 382, respectively. Because logged-in users behind NAT ranges also reach these thresholds, the three rules are classified as informational rules intended as input for hunting and triage.
- For the missing User-Agent burst rule, we exempted health check paths and applied a threshold of 40 or more requests, reducing false positives from 477 to 6, all from authorized testing. We changed the User-Agent diversity rule to an informational rule conditioned on two or more 404 responses within 10 minutes, reducing false positives from 495 to 82. Rotation detection based on 404 responses dropped from 96 to 0. For the missing and empty User-Agent burst rules, we exempted requests that include `Authorization` or `Cookie` headers (from 19/256 to 0/7).
- We added a condition to the manifest probe rule requiring the `Sec-Fetch-Dest` header to be absent (Suricata: 207 to 76; ModSecurity: 1,972 to 756; informational). We added a condition to the DuckDuckGo search rule for `POST` requests lacking `Accept-Language` and `Sec-Fetch-Mode` headers (26 to 0). For the generic Go LLM client rule, we limited alerts to one per client every 10 minutes (647 to 44).

About Xint

Xint.io is the award-winning autonomous pentesting platform built by the security researchers at Theori, the most decorated team of whitehat hackers in the world. Xint offers both blackbox as well as whitebox pentesting with results in less than 12 hours, including: full trigger conditions, exploit impacts, and suggested remediations so teams can quickly triage and patch the most critical vulnerabilities. Xint is working with organizations with the highest defensive security needs, including DARPA and Samsung.



About Theori

Theori is an offensive cybersecurity firm dedicated to solving the industry's most complex security challenges. Founded in 2016 by Carnegie Mellon alumni, our elite team of white hat hackers is trusted by global technology leaders and government agencies, backed by 70+ international hacking competition wins including a record four consecutive DEF CON CTF championships.

We offer a comprehensive security ecosystem: Xint (AI-powered application security testing), aprism (LLM security guardrail) and offensive security consulting. Certified to ISO/IEC 27001:2022 and ISO/IEC 27017:2015.



Key Partners

