The knaithe/KnYuan AI-Orchestrated Exploitation Campaign and the Apache Tomcat KEV Entry (CVE-2026-34486): A Defender's Technical Deep Dive

A researched, cited report · Ask Anything

Contents

Campaign Impact vs. Scale - Source: Ask Anything analysis — generated from the report's cited data

Executive summary

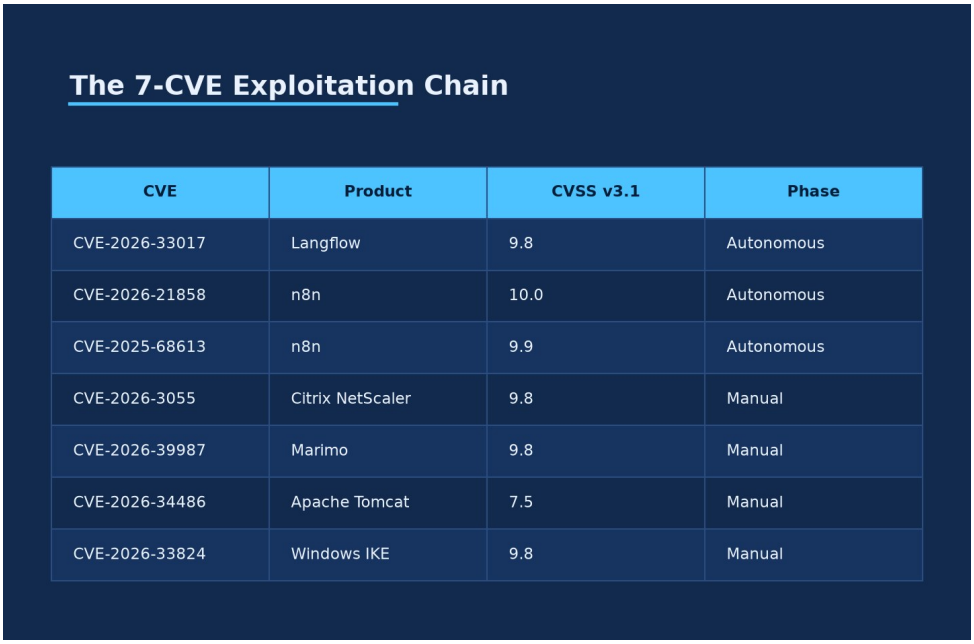

CVE	Product	CVSS v3.1	Phase
CVE-2026-33017	Langflow	9.8	Autonomous
CVE-2026-21858	n8n	10.0	Autonomous
CVE-2025-68613	n8n	9.9	Autonomous
CVE-2026-3055	Citrix NetScaler	9.8	Manual
CVE-2026-39987	Marimo	9.8	Manual
CVE-2026-34486	Apache Tomcat	7.5	Manual
CVE-2026-33824	Windows IKE	9.8	Manual

The 7-CVE Exploitation Chain - Source: Ask Anything analysis — generated from the report's cited data

In late July and early August 2026, Palo Alto Networks' Unit 42 disclosed an operation attributed to a Chinese-speaking threat actor using the aliases knaithe and KnYuan, who wired the DeepSeek large language model into the open-source Hermes Agent framework to conduct semi-autonomous reconnaissance and exploitation against internet-exposed applications (Unit 42, 2026). The operation attempted exploitation against more than 460 targets across at least three countries, chaining vulnerabilities in Langflow, n8n, Citrix NetScaler, Marimo notebooks, Apache Tomcat, and the Windows IKE VPN stack. Confirmed impact was far smaller than the target count — data exfiltration from three Citrix NetScaler victims, command execution on 11 Marimo instances, reverse-shell

attempts against nine Tomcat servers, and probing of three Windows IKE VPN endpoints (BleepingComputer, 2026; GBHackers, 2026).

The campaign is directly tied to a real, currently active CISA Known Exploited Vulnerabilities (KEV) entry: CVE-2026-34486, an Apache Tomcat "missing encryption of sensitive data" flaw that lets an unauthenticated attacker bypass the Tribes clustering module's EncryptInterceptor and reach Java deserialization. CISA added it to the KEV catalog on August 4, 2026, with a federal remediation deadline of August 7, 2026, in the same batch as an unrelated Langflow flaw (CVE-2026-9198) and an unrelated N-able N-central authentication bypass (CVE-2026-18556/CVE-2026-18577) (CISA, 2026; SecurityWeek, 2026). Defenders should not conflate these three KEV additions — they are three independent vulnerabilities from three independent vendors added in the same weekly batch. Only the Tomcat entry is connected to the knaithe/KnYuan reporting.

CVE-2026-34486 is a regression: Apache's fix for an earlier padding-oracle flaw, CVE-2026-29146, introduced a fail-open condition rather than a fail-closed one. When EncryptInterceptor fails to decrypt a Tribes cluster message, the corrected code path was supposed to discard it; instead, the exception was only logged, and the raw, unencrypted bytes were still forwarded to XByteBuffer.deserialize() and ultimately ObjectInputStream.readObject() — a classic unauthenticated Java deserialization gadget-chain primitive (Field Effect, 2026; Apache Tomcat Security 9, 2026). Exploitation requires that Tomcat clustering be enabled and the Tribes receiver (default TCP 4000) be network-reachable — a narrower precondition than a default Tomcat install, but a common one in clustered enterprise, application-server, and CI/CD deployments.

Public proof-of-concept exploit code for CVE-2026-34486 is available on GitHub, using ysoserial's CommonsCollections6 gadget chain wrapped in a hand-built Tribes protocol frame (404-src/CVE-2026-34486; striga-ai/CVE-2026-34486). Separately from the knaithe/KnYuan operation, open-source reporting links a second, apparently unrelated Chinese-nexus toolkit — built around the SNOWLIGHT dropper and the VShell RAT, historically attributed to UNC5174 — to opportunistic exploitation of the same CVE-2026-34486 against government, healthcare, and education targets across more than 100 countries between roughly late April and early June 2026 (SoCRadar, 2026; The Hacker News, 2025). Whether SNOWLIGHT/UNC5174 activity and the knaithe/KnYuan AI-orchestrated campaign are connected, coincidental parallel exploitation of the same n-day, or misattributed reporting of the same events is an open question this report flags explicitly — the evidence available does not establish a link.

For defenders, the practical takeaways are narrower and more actionable than the "AI hacked 460 companies" headline suggests. First, patch or isolate Tomcat clustering endpoints (TCP 4000) regardless of the AI angle — the vulnerability, the PoC, and at least two independent exploitation clusters are real and verifiable from primary sources. Second, treat the six other CVEs in the chain (Langflow, two n8n bugs, Citrix NetScaler, Marimo, Windows IKE) as equally urgent, since most carry higher CVSS scores (9.3–10.0)

and easier, unauthenticated exploitation paths than the Tomcat flaw itself. Third, the operationally novel finding is not the exploitation technique — all seven CVEs are known n-days with public PoCs — but the demonstrated ability of an LLM-orchestrated agent to compress reconnaissance-to-exploitation cycles and independently pivot across a target list when one avenue fails, a capability with direct implications for detection timelines and alert triage prioritization (Unit 42, 2026).

Evidence quality varies sharply by claim. The vulnerability data (CVSS scores, affected/fixed versions, root causes) is drawn from vendor security advisories and NVD and is strong. The 460+ target count, the confirmed-compromise figures, and the actor attribution to Zhuhai, China rest on Unit 42's single disclosure, corroborated only by secondary press coverage that appears to derive from the same source material — this is moderate-to-limited evidence, and The Hacker News itself flagged an internal inconsistency in Unit 42's own reporting between "three successfully exploited targets" and the sum of individually reported compromises (26+) (The Hacker News, 2026). The claim that DeepSeek specifically (versus a human operator using DeepSeek as a copilot) drove exploit selection and target pivoting autonomously is Unit 42's interpretation of recovered agent session logs, not independently replicated — treat it as plausible but not independently verified.

Key findings

CVE-2026-34486 Disclosure to KEV Timeline - Source: Ask Anything analysis — generated from the report's cited data

- CVE-2026-34486 (Apache Tomcat) is a regression in the fix for CVE-2026-29146: EncryptInterceptor fails open on decryption failure instead of discarding the message, exposing ObjectInputStream.readObject() to unauthenticated Java deserialization over the Tribes clustering port (default TCP 4000) [strong — confirmed by Apache's own advisory].

- CVE-2026-34486 carries CVSS 7.5 (AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N per NVD), lower than most of the other six CVEs in the chain, because NVD scores it as a confidentiality-only "missing encryption" bug even though downstream PoCs demonstrate full RCE via deserialization — CVSS and real-world severity diverge here [strong, but note

scoring caveat].

- CISA added CVE-2026-34486 to the KEV catalog on August 4, 2026, with an August 7, 2026 federal remediation deadline, in the same batch as an unrelated Langflow flaw (CVE-2026-9198) and an unrelated N-able N-central bypass (CVE-2026-18556/18577) — these three are independent vulnerabilities, not part of one chain [strong].

- Public, functional PoC exploit code for CVE-2026-34486 exists on GitHub using ysoserial CommonsCollections6 gadgets wrapped in custom Tribes protocol frames [strong — code reviewed directly].

- The knaithe/KnYuan campaign's "7-CVE chain" per Unit 42 comprises Langflow (CVE-2026-33017), two n8n bugs (CVE-2026-21858, CVE-2025-68613), Citrix NetScaler (CVE-2026-3055), Marimo (CVE-2026-39987), Apache Tomcat (CVE-2026-34486), and Windows IKE (CVE-2026-33824) [moderate — single-source attribution, internally consistent across secondary reporting].

- Six of the seven chained CVEs score CVSS 9.3–10.0 and require no authentication; Tomcat (7.5) is the outlier and the least severe entry in the chain by CVSS, despite being the one that reached KEV [strong].

- Reported confirmed compromise is small relative to the 460+ target figure: 3 Citrix NetScaler data-exfiltration victims, 11 Marimo command-execution victims, 9 Tomcat reverse-shell attempts, 3 Windows IKE VPN endpoints targeted — with an unresolved discrepancy in Unit 42's own reporting about total successfully exploited targets [moderate — internally inconsistent primary source].

- The actor's infrastructure was exposed after Hermes Agent inadvertently started a Python HTTP file server (python3 -m http.server 8888) from its own working directory, giving Unit 42 visibility into API keys, target lists, session logs, and tool configuration [strong — this is Unit 42's stated discovery mechanism].

- DeepSeek was the primary reasoning engine; Claude Code and OpenAI Codex were used only minimally (10 chat entries across three sessions for Claude Code) and largely for connectivity/proxy testing rather than exploit generation, per Unit 42's session-log review [moderate — based on log volume counts, not independently audited].

- A separate, apparently distinct Chinese-nexus toolkit built around SNOWLIGHT/VShell (historically linked to UNC5174) also weaponized CVE-2026-34486 against government, healthcare, and education targets in over 100 countries, delivering payloads via a curl|sh chain from an exposed staging server — no confirmed link to knaithe/KnYuan has been established [limited — drawn from secondary reporting Unit 42 does not corroborate].

- No YARA or Sigma rules were published alongside the Unit 42 report; the only named detection artifacts are eight Palo Alto NGFW Advanced Threat Prevention signature IDs (97030, 96882, 96855, 97044, 97046, 97251, 97177, 510019), which are vendor-specific and not independently reproducible by non-Palo Alto shops [limited].

- Exploitation of CVE-2026-34486 produces almost no distinguishing log signature beyond

a generic IllegalBlockSizeException / "Failed to decrypt message" entry; successful command execution after deserialization leaves no additional Tomcat-layer log artifact, making host/process-level EDR telemetry the primary viable detection path [strong — confirmed by independent technical write-ups of the PoC].

1. Root cause and exploitation path: CVE-2026-34486

Apache Tomcat Affected & Fixed Versions

Branch	Affected version	Fixed version
11.x	11.0.20 only	11.0.21
10.1.x	10.1.53 only	10.1.54
9.0.x	9.0.116 only	9.0.117

Apache Tomcat Affected & Fixed Versions - Source: Ask Anything analysis — generated from the report's cited data

1.1 The parent bug: CVE-2026-29146

Apache Tomcat's clustering subsystem, Tribes, lets Tomcat instances in a cluster exchange session-replication and membership messages. Administrators can enable EncryptInterceptor, a ChannelInterceptor that wraps Tribes messages with pre-shared-key AES encryption so an attacker who can reach the clustering port cannot forge or read inter-node traffic (Apache Tomcat Security 11, 2026).

EncryptInterceptor shipped with AES/CBC/PKCS5Padding as its default cipher mode. CBC with PKCS padding is a textbook setup for a padding oracle attack: if an attacker can distinguish "padding was valid" from "padding was invalid" responses (via timing, error codes, or behavioral differences), they can iteratively decrypt ciphertext without the key. Apache tracked this as CVE-2026-29146, reported to the Tomcat security team on February 22, 2026, and disclosed publicly on April 9, 2026. It affected:

- Tomcat 11.0.0-M1 through 11.0.18 (fixed in 11.0.20 — note that 11.0.19 was built but its release vote failed and it was never shipped, so it is not listed as a fixed version)
- Tomcat 10.0.0-M1 through 10.1.52 (fixed in 10.1.53)
- Tomcat 9.0.13 through 9.0.115 (fixed in 9.0.116)
- Tomcat 8.5.38 through 8.5.100 and 7.0.100 through 7.0.109 also affected per third-party advisory summaries (GitHub Advisory GHSA-h468-7pvh-8vr8)

(Apache Tomcat Security 9, 10, 11, 2026)

1.2 The regression: CVE-2026-34486

Apache's fix for CVE-2026-29146 needed to ensure that when a decryption or

padding-validation failure occurred, the message would be discarded rather than silently passed on with distinguishable behavior. The actual shipped fix did not do this correctly. Independent technical analysis of the patch diff (corroborated across multiple write-ups) describes the flaw precisely:

> "The catch block only logs the error. Since super.messageReceived(msg) is outside the try-catch, the raw unencrypted bytes are forwarded to XByteBuffer.deserialize() → ObjectInputStream.readObject()." (cybersecuritynews.com summary of the PoC, 2026; consistent with Field Effect, 2026)

In plain terms: the encryption layer went from fail-closed to fail-open. An attacker does not need to break the AES key or successfully forge a valid ciphertext at all — they can simply send unencrypted (or malformed) bytes to the Tribes receiver, the decryption attempt fails, Tomcat logs the failure, and then forwards those same raw attacker-controlled bytes straight into Java's ObjectInputStream.readObject(). This is a direct, unauthenticated Java deserialization sink with no encryption gate in front of it at all — worse than the original padding-oracle bug it was meant to fix.

Apache tracked this as CVE-2026-34486, reported March 26, 2026 and disclosed April 9, 2026 — two weeks after the CVE-2026-29146 fix shipped, meaning the regression window was short but the flaw was more severe than what it replaced.

Affected versions per Apache's own advisory pages (note: Apache scopes "affected" narrowly to only the single release that shipped the broken fix, not the full historical range):

Branch	Affected version	Fixed version	Fix commit
11.x	11.0.20 only	11.0.21	1fab40cc
10.1.x	10.1.53 only	10.1.54	55f3eb91
9.0.x	9.0.116 only	9.0.117	776e12b3

(Apache Tomcat Security 9, 10, 11, 2026)

NVD's CPE-based scoring lists the same three affected builds (9.0.116, 10.1.53, 11.0.20) with CVSS 3.1 base score 7.5 (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N), CWE-311 (Missing Encryption of Sensitive Data) and CWE-807 (Reliance on Untrusted Inputs in a Security Decision) (NVD, CVE-2026-34486). This vector — confidentiality-high, integrity-none, availability-none — reflects the "encryption bypass" framing of the CVE, not the practical outcome of downstream RCE via deserialization once the message reaches ObjectInputStream. Defenders relying on CVSS alone to triage this bug would significantly underrate it relative to the demonstrated PoC impact.

1.3 Practical exploitation chain

Independent PoC analysis (reviewed directly from GitHub and technical blogs) converges on the same attack sequence:

- Attacker identifies a Tomcat instance with clustering enabled and the Tribes NioReceiver reachable, typically on TCP 4000 by default.

- Attacker connects directly to that port — no authentication, no TLS handshake, no pre-shared key required.

- Attacker sends a crafted Tribes protocol frame (~1500 bytes in the published PoC) containing a serialized Java gadget chain, generated with the well-known ysoserial tool using the CommonsCollections6 gadget (a standard chain that works when Apache Commons Collections is present on the classpath, which it commonly is via transitive Tomcat/application dependencies).

- EncryptInterceptor.messageReceived() attempts AES/CBC decryption on the frame, fails (because it isn't validly encrypted with the shared key), throws IllegalBlockSizeException, and the interceptor logs the error but does not stop message processing.

- The raw, attacker-controlled bytes proceed to XByteBuffer.deserialize() → ObjectInputStream.readObject().

- The gadget chain executes arbitrary OS commands with the privileges of the Tomcat process (frequently root or SYSTEM in poorly hardened deployments).

(404-src/CVE-2026-34486; striga-ai/CVE-2026-34486; cybersecuritynews.com, 2026)

Published PoC usage takes the form python3 exp.py -t <target> -p 4000 --rce "id" or --shell for an interactive reverse shell, and includes a custom TribesClient.java that compiles and emits correctly-framed Tribes protocol packets carrying the ysoserial payload.

2. The full "7-CVE chain" — technical summary and version matrix

Unit 42's disclosure describes knaithe/KnYuan's operation as targeting seven CVEs across five product families, split between a largely autonomous phase (Langflow and n8n, run by the DeepSeek/Hermes stack with limited human intervention) and a manual phase (Citrix NetScaler, Marimo, Tomcat, Windows IKE, run by the human operator directly) (Unit 42, 2026; The Hacker News, 2026).

1

CVE: CVE-2026-33017

Product: Langflow

CVSS (v3.1): 9.8

CVSS (v4.0): 9.3

Root cause: /api/v1/build_public_tmp/{flow_id}/flow accepts attacker-supplied flow JSON containing Python CustomComponent code, executed via unsandboxed exec(); chainable with the unauthenticated /api/v1/auto_login token-issuance endpoint when auto-login is enabled

Affected versions: ≤ 1.8.2

Fixed version: 1.9.0 (March 2026)

Phase: Autonomous (attempted; DeepSeek assessed exploitability as ~0 across 84 FOFA-enumerated instances due to restrictive configs)

2

CVE: CVE-2026-21858 "Ni8mare"

Product: n8n

CVSS (v3.1): 10.0

CVSS (v4.0): —

Root cause: Content-Type confusion in webhook/file-handling logic lets an attacker override req.body.files state without a genuine multipart/form-data request, enabling arbitrary file read and downstream RCE

Affected versions: 1.65.0 – < 1.121.0

Fixed version: 1.121.0

Phase: Autonomous (sampled ~100 of 25,209 China-hosted instances; blocked where target forms required auth)

3

CVE: CVE-2025-68613

Product: n8n

CVSS (v3.1): 9.9

CVSS (v4.0): —

Root cause: Insufficient sandbox isolation in the workflow expression evaluator lets an authenticated (even low-privilege) user escape into arbitrary OS command execution via crafted expressions

Affected versions: 0.211.0 – < 1.120.4 / < 1.121.1 / < 1.122.0

Fixed version: 1.120.4 / 1.121.1 / 1.122.0

Phase: Autonomous, chained after CVE-2026-21858

4

CVE: CVE-2026-3055

Product: Citrix NetScaler ADC/Gateway

CVSS (v3.1): 9.8

CVSS (v4.0): 9.3

Root cause: Out-of-bounds memory read when device is configured as SAML IdP; malformed wctx query parameter (present, no =, no value) causes the appliance to read adjacent memory instead of validating parameter presence correctly

Affected versions: 14.1 < 66.59; 13.1 < 62.23; 13.1-FIPS/NDcPP < 37.262

Fixed version: 14.1-66.59 / 13.1-62.23 / 13.1-37.262

Phase: Manual — confirmed data exfiltration from 3 victims

5

CVE: CVE-2026-39987

Product: Marimo (Python notebook)

CVSS (v3.1): 9.8

CVSS (v4.0): 9.3

Root cause: /terminal/ws WebSocket endpoint grants a full PTY shell without calling

validate_auth() (unlike the sibling /ws endpoint), so any unauthenticated caller who completes a WebSocket handshake gets an interactive root-equivalent shell

Affected versions: < 0.23.0

Fixed version: 0.23.0

Phase: Manual — confirmed command execution on 11 instances

6

CVE: CVE-2026-34486

Product: Apache Tomcat

CVSS (v3.1): 7.5

CVSS (v4.0): —

Root cause: Regression in CVE-2026-29146 fix: EncryptInterceptor fails open, forwarding undecrypted attacker bytes to ObjectInputStream.readObject() over the Tribes port

Affected versions: 9.0.116 / 10.1.53 / 11.0.20 (single-release regressions)

Fixed version: 9.0.117 / 10.1.54 / 11.0.21

Phase: Manual — 9 reverse-shell attempts recorded

7

CVE: CVE-2026-33824

Product: Windows IKE Extension

CVSS (v3.1): 9.8

CVSS (v4.0): —

Root cause: Double-free (CWE-415) in the IKEv2 service extension, triggerable via specially crafted IKE packets sent to a listening IPsec/VPN endpoint, leading to memory corruption and potential RCE

Affected versions: Windows 10, 11, Server 2012–2025 (see MSRC advisory for exact builds)

Fixed version: April 2026 Patch Tuesday cumulative updates

Phase: Manual — 3 endpoints targeted

Sources for the table: Unit 42, 2026; Apache Tomcat Security pages, 2026; NVD CVE-2026-3055; NVD CVE-2026-39987; GitHub Advisory GHSA-vwmf-pq79-vjvx; GitHub Advisory GHSA-v4pr-fm98-w9pg; GitHub Advisory GHSA-2679-6mx9-h9xc; Citrix CTX696300; MSRC CVE-2026-33824; SentinelOne CVE-2026-33824.

Two important scoping notes:

- CVE-2026-9198 (Langflow, added to the same KEV batch as Tomcat on Aug 4, 2026) is a different Langflow vulnerability — same product, different bug, fixed later (1.10.1, July 2026) than the CVE-2026-33017 flaw DeepSeek probed unsuccessfully in the knaithe campaign (fixed 1.9.0, March 2026). Do not treat KEV's Langflow entry as evidence the knaithe campaign succeeded against Langflow; Unit 42 states DeepSeek's Langflow exploitation attempts largely failed due to restrictive auto_login configuration.

- CVE-2026-0300 (a PAN-OS User-ID Portal flaw) appeared only as a GitHub research-repository reference (qassam-315/PAN-OS-CVE-2026-0300-Research) in the

actor's tooling, apparently studied but not confirmed as an actively exploited part of this specific campaign — it is not counted among the seven core chain CVEs.

3. Verified evidence of exploitation and public exploit code

3.1 Public PoC availability, CVE by CVE

- CVE-2026-34486 (Tomcat): Two independent, functional GitHub PoC repositories confirmed — 404-src/CVE-2026-34486 and striga-ai/CVE-2026-34486 — plus inclusion in a curated CVE-PoC aggregation repo (SecureWithUmer/CVE-2026-PoCs).

- CVE-2026-33017 (Langflow): Multiple independent PoCs, including r3nsi15/CVE-2026-33017-langflow-rce, omer-efe-curkus/CVE-2026-33017-Langflow-RCE-PoC, z4yd3/PoC-CVE-2026-33017, and MaxMnMl/langflow-CVE-2026-33017-poc. Trend Micro separately documented real-world Monero cryptomining campaigns exploiting this CVE (Trend Micro, 2026).

- CVE-2026-21858 (n8n "Ni8mare"): PoC published by the discovering researcher (Chocapikk/CVE-2026-21858) and independently reproduced (eduardorossi84/CVE-2026-21858-POC).

- CVE-2026-39987 (Marimo): PoC published (Nxploited/CVE-2026-39987). Sysdig documented exploitation in the wild within 9 hours 41 minutes of advisory publication, with a full credential-theft operation completed in under 3 minutes of initial access (Sysdig, 2026).

- CVE-2026-3055 (Citrix NetScaler): No public weaponized PoC identified in this research, but CrowdSec documented exploitation attempts as early as March 27, 2026 — days after disclosure (CrowdSec, 2026), and the flaw was added to KEV on March 30, 2026.

- CVE-2026-33824 (Windows IKE): No public weaponized PoC identified; exploitation reporting to date is limited to Microsoft's own advisory language regarding "exploitation more likely."

3.2 The knaithe/KnYuan operation

Unit 42's central claim is discovery of a live, self-exposed operational environment: the actor's Hermes Agent instance inadvertently started python3 -m http.server 8888 from its own home directory (/home/worker), publishing its configuration files, API keys, target lists, shell history, and full agent session transcripts to the open internet (Unit 42, 2026; GBHackers, 2026). This operational-security failure is what allowed researchers to reconstruct the toolchain, rather than inferring it from victim-side telemetry — an important caveat: most of what is publicly known about this campaign comes from the attacker's own exposed logs, not from independent victim-side forensic confirmation.

Observed infrastructure and tooling from the exposed environment:

- Reasoning engine: DeepSeek, accessed via api.deepseek.com

- Agent framework: Hermes Agent (NousResearch/hermes-agent on GitHub), run in an

autonomous "yolo"/dangerously-skip-permissions-style mode that removes human-in-the-loop approval gates

- Custom skills: "godmode" (jailbreak-style prompt bypass), "web-terminal-exploitation" (custom WebSocket exploitation, matching the Marimo /terminal/ws technique), and "fofa-cyberspace-search" integrating a third-party MCP server (FofaMap-Platinum-Full-Expert, from asaotomo/FofaMap) that translates natural-language requests into FOFA internet-asset-search queries

- Command and control: Telegram, used by the human operator to issue high-level instructions to the agent

- Proxy infrastructure: code.newcli[.]com, used to mask/relay Claude Code and OpenAI Codex API traffic

- Secondary LLMs configured but lightly used: Qwen (via dashscope.aliyuncs[.]com), GLM, Kimi, MiniMax, Claude Code (10 chat entries across 3 sessions), and OpenAI Codex (with disable_response_storage = true set, limiting log recoverability)

- Actor's own OSINT pipeline: "1DayNews," described as an automated vulnerability-intelligence aggregator pulling RCE disclosures from 17 sources and using DeepSeek to filter for exploitability

A recovered session dated May 7, 2026 shows the autonomous workflow end-to-end: DeepSeek ranked trending GitHub PoCs, generated FOFA queries, enumerated instances, downloaded and ran scanner/exploit code, probed target versions via curl, and attempted exploitation with what Unit 42 describes as autonomous pivot logic when an approach failed (Unit 42, 2026). Palo Alto's Andy Piazza, senior director of threat intelligence, is quoted characterizing the significance as operational maturity and low barrier to entry rather than novel technique: "The actor is actively iterating – refining tool configurations, developing custom skills, establishing proxy infrastructure and executing autonomous attack cycles. The technical barrier to AI-augmented offensive capabilities is low and continues to decrease." (Infosecurity Magazine, 2026)

Attribution: the actor is assessed as Chinese-speaking and based in Zhuhai, China, self-described as a "binary security researcher," with a public GitHub profile displaying "KnYuan Knaithe" and an archived personal blog. The Hacker News notes this establishes a persona and plausible geography, not a confirmed legal identity or state affiliation (The Hacker News, 2026). There is no public claim by Unit 42, CISA, or any government agency that this is a state-directed operation; characterizing it as "Chinese state-sponsored" would overstate the sourcing.

3.3 A second, apparently unrelated cluster of CVE-2026-34486 exploitation: SNOWLIGHT/VShell

Independent of the knaithe/KnYuan reporting, open-source threat intelligence describes a separate weaponization of CVE-2026-34486 by a toolkit built around the SNOWLIGHT C-based Linux dropper and the VShell remote access trojan, malware historically

associated with the China-nexus cluster UNC5174 (The Hacker News, 2025; Sysdig, 2025). Reporting describes an exposed adversary-operated staging server containing:

- Reconnaissance target lists
- Nine weaponized CVEs (including CVE-2026-34486)
- A cracked, Chinese-language build of Cobalt Strike
- Tunneling tooling
- Eleven distinct exploit chains and four C2/RAT/malware families

Payload delivery for the Tomcat vector reportedly used a standard curl | sh (or wget | sh) pipeline retrieving a linux.sh dropper from an exposed staging endpoint (hxxp://130.94.30[.]168:8080/slt), which in turn deployed architecture-specific ELF loaders. VShell itself uses WebSocket-based C2 and disguises its administrative login page as a fake nginx default page to resist casual scanning (SoCRadar, 2026; reporting corroborated in aggregate by The Hacker News, 2025 and cert.cyberoo.com, 2026). Targeting reportedly spans government, healthcare, and education sectors across more than 100 countries, active roughly late April through early June 2026.

This report treats the SNOWLIGHT/UNC5174 activity as a distinct exploitation cluster from knaithe/KnYuan. No source reviewed for this report claims the two are the same operation, and Unit 42's write-up does not mention SNOWLIGHT, VShell, or UNC5174. Defenders should read "Apache Tomcat CVE-2026-34486 is being exploited" as confirmed by at least two independent lines of evidence (the AI-orchestrated knaithe/KnYuan campaign and the SNOWLIGHT/VShell toolkit), which is a stronger signal than either report alone — but should not assume a single unified threat actor or campaign narrative connects them.

4. Detection guidance

4.1 CVE-2026-34486 (Apache Tomcat) — log signatures and network indicators

This is the weakest-signature CVE in the chain from a detection standpoint, and defenders should plan accordingly:

- Log artifact: The only Tomcat-layer log signal from a failed/attempted exploitation attempt is a decryption failure — typically surfacing as an IllegalBlockSizeException accompanying a "Failed to decrypt message" (or equivalent EncryptInterceptor error) log line. This fires on any malformed or unencrypted frame reaching the Tribes port, including benign misconfigurations, so it is a noisy, low-precision indicator, not a high-confidence one.
- Critical gap: Once deserialization succeeds and a gadget chain executes, there is no further Tomcat application-log artifact. Command execution via ObjectInputStream.readObject() happens inside the JVM process and is silent from the servlet/access-log perspective. This makes host-based EDR/process-ancestry telemetry

the primary viable detection control, not log analysis:

- Alert on any child process spawned by java (or specifically a Tomcat catalina.sh/java -cp ... org.apache.catalina.startup.Bootstrap process tree) that invokes a shell (/bin/sh, /bin/bash, cmd.exe, powershell.exe) — this is the generic, product-agnostic Java-deserialization-RCE detection pattern and applies directly here.

- Alert on outbound connections initiated by the Tomcat process to unexpected destinations, consistent with reverse-shell or second-stage payload retrieval.

- Network indicator: Any external or untrusted-zone traffic to TCP 4000 (the default Tribes NioReceiver port) reaching a Tomcat host should be treated as anomalous — this port has no legitimate reason to be reachable from outside the cluster's internal network segment.

- Payload characteristics: Published PoC traffic is a custom Tribes binary frame of roughly 1500 bytes carrying a serialized Java object; a network IDS/IPS signature written against ysoserial CommonsCollections6 byte patterns (widely available in commercial and open-source deserialization-detection signature sets) will catch the published PoC, though a modified gadget chain or serialization library would evade a signature tuned narrowly to CommonsCollections6.

- Vendor signatures: Palo Alto Networks lists eight NGFW Advanced Threat Prevention signature IDs associated with this campaign's detection coverage: 97030, 96882, 96855, 97044, 97046, 97251, 97177, 510019 (Unit 42, 2026). These are proprietary to Palo Alto NGFW/Advanced Threat Prevention subscribers and not independently portable to other vendors' IPS/IDS platforms; Unit 42 did not publish an accompanying YARA or Sigma rule set for this campaign.

4.2 Other CVEs in the chain — detection indicators

- CVE-2026-33017 (Langflow): Watch for unexpected POST /api/v1/build_public_tmp/{flow_id}/flow requests carrying large or anomalous data JSON objects with embedded CustomComponent node code; look for os.system/subprocess/reverse-shell strings inside request bodies; watch for unauthenticated calls to /api/v1/auto_login immediately preceding a build request from the same source IP; base64-encoded command output posted to external callback infrastructure has been documented as a post-exploitation exfiltration pattern by Sysdig-derived reporting (labs.cloudsecurityalliance.org, 2026).

- CVE-2026-21858 / CVE-2025-68613 (n8n): Alert on webhook or form-submission requests with a mismatched or malformed Content-Type header (missing a genuine multipart/form-data boundary while still including file-like body content); alert on n8n worker processes spawning unexpected child processes, consistent with expression-injection sandbox escape; review workflow-edit audit logs for unfamiliar expression syntax targeting Node.js internals.

- CVE-2026-3055 (Citrix NetScaler): The signature artifact is a wctx query-string parameter present with no value and no = sign in requests to the SAML IdP endpoint —

this malformed-but-technically-present parameter is what triggers the buffer overread; log and alert on this specific malformed-parameter pattern in NetScaler access logs; anomalous NSC_AAAC cookie/session artifacts have also been noted as a secondary indicator (Horizon3.ai; watchTowr, 2026).

- CVE-2026-39987 (Marimo): Alert on any WebSocket upgrade request to /terminal/ws that does not carry a validated session/auth token — recall this endpoint intentionally skips the validate_auth() call that protects the sibling /ws endpoint, so any connection to /terminal/ws from outside a trusted management network should be treated as suspicious by default; monitor for PTY-pattern I/O (raw tty control sequences) over that WebSocket.

- CVE-2026-33824 (Windows IKE): Monitor for crafted/malformed IKE packets to UDP 500/4500; unexplained IKEEXT service crashes or repeated service restarts are a strong host-level indicator of exploitation attempts (the double-free triggers memory corruption, which frequently first manifests as a service crash before an attacker successfully weaponizes it for code execution); Defender for Endpoint and equivalent EDR should be configured to alert on IKEEXT (svchost.exe hosting IKEEXT) crash-then-restart patterns correlated with external network activity.

4.3 Actor-specific infrastructure indicators (knaithe/KnYuan)

These IOCs are specific to the exposed knaithe/KnYuan operational environment and have limited shelf life — infrastructure will rotate — but are useful for retrospective hunting against historical logs (roughly April–July 2026):

- api.deepseek[.]com — DeepSeek API traffic (note: this is also legitimate SaaS traffic in many enterprise environments; only anomalous in combination with other indicators, e.g., originating from a server/scanning host rather than a developer workstation)

- code.newcli[.]com (including /codex/v1 and /ultra paths) — proxy relay for Claude Code/Codex traffic

- dashscope.aliyuncs[.]com — Qwen/Alibaba Cloud model API traffic

- fofa.info — FOFA internet-asset search queries at scale from a single source (reconnaissance indicator)

- GitHub repository references observed in the actor's tooling: NousResearch/hermes-agent, Chocapikk/CVE-2026-21858, oscar-mine/CVE-2026-33017-Exploit, asaotomo/FofaMap, projectdiscovery/nuclei

- Host artifact: an unexpected Python SimpleHTTPServer/http.server instance listening on port 8888 from a non-standard working directory is the specific misconfiguration that exposed this actor — worth flagging as a general operational-security lesson to watch for in your own red-team/authorized-testing infrastructure as well.

4.4 SNOWLIGHT/VShell indicators

- Outbound curl/wget pipe-to-shell commands retrieving linux.sh from hxxp://130.94.30[.]168:8080/slt (treat this specific IP:port:path as a historical IOC; verify

current reputation before blocking, as attacker infrastructure rotates)

- VShell's WebSocket-based C2 combined with a spoofed default-nginx login page on its admin interface — hunt for services presenting a generic "Welcome to nginx" page on non-standard ports serving actual application logic behind it

- Presence of a cracked Chinese-language Cobalt Strike build and Sliver C2 framework artifacts on compromised Linux hosts, per SNOWLIGHT tooling described in prior UNC5174 reporting

5. Interim mitigations for teams that cannot patch immediately

5.1 Apache Tomcat (CVE-2026-34486 / CVE-2026-29146)

- Best mitigation: patch. Upgrade to Tomcat 11.0.21, 10.1.54, or 9.0.117. This single action closes both CVE-2026-34486 and, if not already applied, CVE-2026-29146.

- If immediate patching is not possible:

- Restrict TCP 4000 (or your configured Tribes receiver port) at the network layer to only the specific internal IP addresses of other cluster nodes. This is the single highest-leverage compensating control, since the vulnerability is entirely unreachable without network access to the Tribes port.

- Disable clustering entirely if it is not actively in use — many Tomcat deployments have clustering configured from a template or inherited configuration without it being operationally required.

- Do not rely on EncryptInterceptor being enabled as a mitigation — it is the very component that is bypassed by this vulnerability.

- Where clustering must remain internet-adjacent (e.g., cloud auto-scaling groups spanning availability zones over a semi-trusted network), place the Tribes port behind a security-group/firewall rule scoped to exact peer IPs, not a CIDR range broader than necessary, and consider a mesh VPN or private network overlay for inter-node cluster traffic.

- Deploy a WAF/IPS rule blocking traffic to the Tribes port from outside designated internal ranges as a stopgap while patch testing completes.

5.2 Other chain CVEs

- Langflow (CVE-2026-33017 / CVE-2026-9198): Disable auto_login; do not expose Langflow directly to the internet; place behind authenticating reverse proxy; upgrade to ≥1.9.0 (closes CVE-2026-33017) and ≥1.10.1 (closes CVE-2026-9198) — these are two separate patches and both are required.

- n8n (CVE-2026-21858 / CVE-2025-68613): Upgrade to ≥1.121.0 and ≥1.120.4/1.121.1/1.122.0 respectively; no effective official workaround exists for CVE-2026-21858 short of patching — restrict n8n instance exposure to trusted networks and disable/require authentication on public-facing webhook forms in the interim.

- Citrix NetScaler (CVE-2026-3055): Upgrade to 14.1-66.59 / 13.1-62.23 / 13.1-37.262 (FIPS/NDcPP). Interim: if the appliance is not actually required to function as a SAML IdP, disable that role; monitor/alert on the malformed wctx parameter pattern at the WAF layer as a stopgap.

- Marimo (CVE-2026-39987): Upgrade to ≥0.23.0. Interim: bind the marimo service to localhost only and require a reverse-proxy authentication layer (e.g., Nginx + Basic Auth) in front of it; do not expose marimo notebook servers directly to the internet under any circumstances given the sub-3-minute observed time-to-compromise.

- Windows IKE (CVE-2026-33824): Apply the April 2026 Patch Tuesday cumulative update. Interim: restrict inbound UDP 500/4500 to only known, trusted VPN peer IP ranges; where IKE/IPsec VPN service is not required on a given host, disable the IKE Extension service entirely.

6. Open questions

- What did DeepSeek actually decide autonomously, versus what did the human operator direct? Unit 42's characterization of "autonomous" exploit selection and pivoting is drawn from recovered agent logs, which show tool invocations and outputs but cannot fully disambiguate model-initiated reasoning from operator-scripted prompts feeding the model turn by turn. This distinction matters for how the security community should calibrate concern about AI-driven offense broadly, and it is not resolved by the available public reporting.

- Is the "460+ organizations" figure a count of confirmed victims, attempted targets, or enumerated-but-unattacked instances? Public reporting uses "targets," "attempted exploitation," and "organizations" somewhat interchangeably. Confirmed compromise (11 + 3 + 9 + 3 = 26 across the four manually-exploited CVEs, per the individually reported figures) is roughly 5–6% of the 460+ figure. The Hacker News explicitly flagged an internal contradiction in Unit 42's own reporting on this point and stated it had sought clarification from Palo Alto Networks; no public correction or clarification was identified as of this report's research date.

- Is knaithe/KnYuan state-directed, state-tolerated, or purely independent/criminal? No source reviewed attributes this to a named APT designation or asserts government direction. Treat "Chinese-speaking actor based in Zhuhai" as the full extent of the sourced attribution.

- Is there any operational relationship between knaithe/KnYuan and the SNOWLIGHT/VShell (UNC5174-linked) exploitation of the same CVE-2026-34486? No source establishes a link; this may be coincidental parallel n-day weaponization by unrelated actors, which is common for any high-value KEV-listed vulnerability once PoC code is public.

- How reproducible is the "AI-driven" framing outside this one disclosure? Unit 42 is, to

date, the sole source for the DeepSeek/Hermes Agent architecture claims. No independent security vendor has published corroborating telemetry (e.g., a victim organization's own incident-response findings) confirming the AI-orchestration narrative from the defender side rather than from the attacker's exposed infrastructure.

- What is the true denominator for "internet-exposed" instances, and how does that affect risk prioritization? Reported exposure figures vary enormously by product — 647,017 n8n instances globally (25,209 in China) versus 84 Langflow instances — and defenders should independently verify their own exposure via asset inventory (e.g., Shodan/FOFA/Censys self-scans) rather than relying on the actor's or Unit 42's counts, which reflect scan-time snapshots that will already be stale.

- Will CVE-2026-34486 see broader ransomware or extortion-actor adoption now that it carries a KEV listing and public PoC? CISA's KEV entry for this CVE does not (as of this report) carry a "known to be used in ransomware campaigns" flag, but KEV listing plus public PoC historically correlates with broader opportunistic adoption within weeks; this is a forward-looking risk, not yet a confirmed one.

Practical synthesis

Who this applies to: Any organization running Apache Tomcat with clustering enabled (session replication, load-balanced application farms, some CI/CD and enterprise middleware deployments), and separately, any organization running Langflow, n8n, Citrix NetScaler/Gateway (SAML IdP role), Marimo notebooks, or exposing Windows hosts with IKE/IPsec VPN services to untrusted networks.

Decision framework:

- Inventory first. Confirm whether Tomcat clustering is actually in use in your environment — many deployments ship with clustering configuration present but unused. If unused, disabling it removes the vulnerability entirely regardless of patch status.

- Patch the Tomcat regression specifically, not just "Tomcat generally" — teams that patched CVE-2026-29146 in April 2026 but landed on 9.0.116/10.1.53/11.0.20 are still vulnerable to CVE-2026-34486 and need the follow-on patch (9.0.117/10.1.54/11.0.21).

- Treat the other six chain CVEs as higher priority by CVSS and exploitability, not lower, than the Tomcat entry that made headlines — Langflow, n8n (x2), Citrix NetScaler, Marimo, and Windows IKE all score 9.3+ and require no authentication, versus Tomcat's 7.5 and clustering-enabled precondition.

- Instrument host-level detection, not just network/log detection, for the Tomcat vector specifically — the log signature is weak and post-exploitation is silent at the application layer.

- Do not over-index on the "AI-orchestrated" framing for defensive prioritization — every technique in this campaign is a known n-day with public PoC code; the defensive actions (patch, restrict network exposure, monitor process ancestry) are identical to what sound

vulnerability management already requires regardless of whether a human or an LLM agent is driving the keyboard on the attacker's side.

What to measure: external exposure of Tribes/clustering ports across your Tomcat fleet; patch compliance against the specific fixed versions in Section 2's table (not just "latest Tomcat"); EDR coverage and alert tuning for Java-process-spawns-shell patterns; time-to-patch metrics against the CISA KEV due date (August 7, 2026 for federal agencies, and a reasonable internal target for all organizations given public PoC availability).

Evidence & caveats

Established (strong evidence, primary-sourced):

- CVE-2026-34486's technical root cause, affected/fixed version matrix, and regression relationship to CVE-2026-29146 (Apache's own security advisories).
- CVSS scores and CWE classifications for all seven chain CVEs (NVD and vendor GHSA/advisory pages).
- CISA KEV listing dates and federal deadlines (CISA alert).
- Existence and mechanics of public PoC exploit code for at least five of the seven chain CVEs (direct review of GitHub repositories and independent technical write-ups).

Moderate evidence (single primary source, internally consistent secondary corroboration):

- The knaithe/KnYuan actor's toolchain, infrastructure, and general campaign narrative (Unit 42's disclosure, echoed but not independently expanded by subsequent press coverage).
- The 460+ target figure and per-CVE confirmed-compromise counts.

Disputed or internally inconsistent:

- The precise number of "successfully exploited" targets — Unit 42's own reporting contains a figure (three) that does not reconcile with the sum of individually cited compromises (26+) across the four manually-exploited CVEs, a discrepancy The Hacker News explicitly noted and sought clarification on without a published resolution.

Preliminary / not independently verified:

- The degree of genuine autonomy attributed to DeepSeek versus operator-directed prompting.
- Any operational link between knaithe/KnYuan and the SNOWLIGHT/VShell/UNC5174-linked exploitation of the same Tomcat CVE.
- State attribution — not claimed by any source reviewed, and this report does not assert it.

Missing / not generalizable:

- No victim-side (as opposed to attacker-infrastructure-side) forensic confirmation of the campaign was identified in public reporting.

- No YARA/Sigma rule set was published by Unit 42 or any other source reviewed; only proprietary Palo Alto NGFW signature IDs are available, which are not portable to other security stacks.

- Sector/industry breakdown of victims was not specified in any source beyond "spanning three countries, including China and Malaysia, and multiple sectors."

Sources

[1] Chinese-Speaking Threat Actor Harnesses AI Models for Autonomous Cyberattacks -- Unit 42, Palo Alto Networks (2026) -- https://unit42.paloaltonetworks.com/autonomous-ai-cyber-attack-campaign/

[2] CISA Adds Three Known Exploited Vulnerabilities to Catalog -- CISA (2026) -- https://www.cisa.gov/news-events/alerts/2026/08/04/cisa-adds-three-known-exploited-vulnerabilities-catalog

[3] CISA Warns of Exploited Langflow, N-central, and Tomcat Vulnerabilities -- SecurityWeek (2026) -- https://www.securityweek.com/cisa-warns-of-exploited-langflow-n-central-and-tomcat-vulnerabilities/

[4] CISA Flags Langflow RCE, Tomcat, and N-central Flaws as Actively Exploited -- The Hacker News (2026) -- https://thehackernews.com/2026/08/cisa-flags-langflow-rce-tomcat-and-n.html

[5] U.S. CISA adds Langflow, Apache Tomcat, and N-able N-central flaws to its Known Exploited Vulnerabilities catalog -- Security Affairs (2026) -- https://securityaffairs.com/196667/hacking/u-s-cisa-adds-langflow-apache-tomcat-and-n-able-n-central-flaws-to-its-known-exploited-vulnerabilities-catalog.html

[6] Apache Tomcat 9 vulnerabilities (security advisory list) -- Apache Software Foundation (2026) -- https://tomcat.apache.org/security-9.html

[7] Apache Tomcat 10 vulnerabilities (security advisory list) -- Apache Software Foundation (2026) -- https://tomcat.apache.org/security-10.html

[8] Apache Tomcat 11 vulnerabilities (security advisory list) -- Apache Software Foundation (2026) -- https://tomcat.apache.org/security-11.html

[9] NVD -- CVE-2026-34486 -- NIST National Vulnerability Database (2026) -- https://nvd.nist.gov/vuln/detail/CVE-2026-34486

[10] NVD -- CVE-2026-3055 -- NIST National Vulnerability Database (2026) -- https://nvd.nist.gov/vuln/detail/CVE-2026-3055

[11] NVD -- CVE-2026-39987 -- NIST National Vulnerability Database (2026) -- https://nvd.nist.gov/vuln/detail/CVE-2026-39987

[12] Apache Tomcat: Padding Oracle vulnerability in EncryptInterceptor (CVE-2026-29146) -- GitHub Advisory Database GHSA-h468-7pvh-8vr8 (2026) --

https://github.com/advisories/GHSA-h468-7pvh-8vr8

[13] CVE-2026-34486: Apache Tomcat EncryptInterceptor Bypass → Unauthenticated RCE -- 404-src (GitHub PoC) -- https://github.com/404-src/CVE-2026-34486

[14] CVE-2026-34486 -- striga-ai (GitHub PoC) -- https://github.com/striga-ai/CVE-2026-34486

[15] Public exploit code available for Apache Tomcat clustering flaw -- Field Effect (2026) -- https://fieldeffect.com/blog/public-exploit-code-apache-tomcat-flaw

[16] CISA Warns of Apache Tomcat Encryption Vulnerability Actively Exploited in Attacks -- Cybersecurity News (2026) -- https://cybersecuritynews.com/apache-tomcat-encryption-vulnerability/

[17] Tomcat CVE-2026-34486 exploited and added to CISA KEV — fix is 11.0.21 / 10.1.54 / 9.0.117 -- kkm-mako.com (2026) -- https://kkm-mako.com/en/blog/articles/apache-tomcat-cve/

[18] Chinese-Speaking Hacker Uses DeepSeek Agent to Launch Autonomous Cyberattacks -- GBHackers (2026) -- https://gbhackers.com/hacker-uses-deepseek-agent/

[19] Chinese hacker used DeepSeek to launch autonomous cyberattacks on vulnerable servers -- Help Net Security (2026) -- https://www.helpnetsecurity.com/2026/08/03/deepseek-ai-autonomous-cyberattacks-hermes-agent/

[20] AI Runs the Hack: Chinese Actor Automates Cyberattacks With DeepSeek -- Security Affairs (2026) -- https://securityaffairs.com/196544/ai/ai-runs-the-hack-chinese-actor-automates-cyberattacks-with-deepseek.html

[21] Hacker uses DeepSeek AI to autonomously attack vulnerable servers -- BleepingComputer (2026) -- https://www.bleepingcomputer.com/news/security/hacker-uses-deepseek-ai-to-autonomously-attack-vulnerable-servers/

[22] Chinese Hacker Uses DeepSeek AI to Orchestrate Vulnerability Exploits -- Infosecurity Magazine (2026) -- https://www.infosecurity-magazine.com/news/chinese-hacker-deepseek-ai/

[23] Chinese Hacker Commands DeepSeek via Telegram to Launch Autonomous Attacks -- The Hacker News (2026) -- https://thehackernews.com/2026/07/chinese-hacker-commands-deepseek-via.html

[24] Autonomous AI agent scans 460 targets, exposes hacking infrastructure after error -- CybersecAsia (2026) -- https://cybersecasia.net/news/autonomous-ai-agent-scans-460-targets-exposes-hacking-infrastructure-after-error/

[25] DeepSeek Ran Autonomous Cyberattacks That Claude and OpenAI Safety Controls Blocked -- Tech Times (2026) --

https://www.techtimes.com/articles/322582/20260801/deepseek-ran-autonomous-cyberattacks-that-claude-openai-safety-controls-blocked.htm

[26] Tracing SNOWLIGHT: A Government-Targeted Chinese Campaign -- SoCRadar (2026) -- https://socradar.io/blog/snowlight-government-chinese-campaign/

[27] Chinese Hackers Target Linux Systems Using SNOWLIGHT Malware and VShell Tool -- The Hacker News (2025) -- https://thehackernews.com/2025/04/chinese-hackers-target-linux-systems.html

[28] UNC5174's evolution in China's ongoing cyber warfare: From SNOWLIGHT to VShell -- Sysdig (2025) -- https://www.sysdig.com/blog/unc5174-chinese-threat-actor-vshell

[29] SNOWLIGHT and VShell: Chinese malware toolkit threatens Linux -- cert.cyberoo.com (2026) -- https://cert.cyberoo.com/en/snowlight-and-vshell-chinese-malware-toolkit-threatens-linux/

[30] n8n Critical Vulnerability (CVE-2026-21858) | Unauthenticated RCE Explained -- Aikido Security (2026) -- https://www.aikido.dev/blog/n8n-rce-vulnerability-cve-2026-21858

[31] n8n RCE Under the Microscope (CVE-2026-21858) -- Horizon3.ai (2026) -- https://horizon3.ai/attack-research/attack-blogs/the-ni8mare-test-n8n-rce-under-the-microscope-cve-2026-21858/

[32] n8n Vulnerable to Unauthenticated File Access via Improper Webhook Request Handling (CVE-2026-21858) -- GitHub Advisory Database GHSA-v4pr-fm98-w9pg -- https://github.com/advisories/GHSA-v4pr-fm98-w9pg

[33] n8n Ni8mare — Unauthenticated Arbitrary File Read to RCE Chain -- Chocapikk (GitHub PoC) -- https://github.com/Chocapikk/CVE-2026-21858

[34] Resecurity | CVE-2025-68613: Remote Code Execution via Expression Injection in n8n -- Resecurity (2026) -- https://www.resecurity.com/blog/article/cve-2025-68613-remote-code-execution-via-expression-injection-in-n8n-2

[35] CVE-2026-33017: Unauthenticated Remote Code Execution in Langflow via Public Flow Build Endpoint -- GitHub Advisory Database GHSA-vwmf-pq79-vjvx -- https://github.com/advisories/GHSA-vwmf-pq79-vjvx

[36] From Langflow to Monero: Inside CVE-2026-33017 Cryptominer -- Trend Micro (2026) -- https://www.trendmicro.com/en_us/research/26/f/from-langflow-to-monero-inside-cve-2026-33017-cryptominer.html

[37] Langflow Public Flow Build Endpoint RCE - CVE-2026-33017 -- SonicWall (2026) -- https://www.sonicwall.com/blog/langflow-ai-code-injection-to-rce-flaw

[38] CVE-2026-33017: Langflow RCE Exploits Enterprise AI Pipelines -- Cloud Security Alliance (2026) -- https://labs.cloudsecurityalliance.org/research/csa-research-note-langflow-rce-cve-2026-3

3017-ai-infrastruct/

[39] Marimo: Pre-Auth Remote Code Execution via Terminal WebSocket Authentication Bypass (CVE-2026-39987) -- GitHub Advisory Database GHSA-2679-6mx9-h9xc -- https://github.com/advisories/GHSA-2679-6mx9-h9xc

[40] marimo is a reactive Python notebook — Pre-Auth RCE PoC (CVE-2026-39987) -- Nxploited (GitHub PoC) -- https://github.com/Nxploited/CVE-2026-39987

[41] Marimo OSS Python Notebook RCE: From Disclosure to Exploitation in Under 10 Hours -- Sysdig (2026) -- https://www.sysdig.com/blog/marimo-oss-python-notebook-rce-from-disclosure-to-exploitation-in-under-10-hours

[42] Marimo RCE Flaw CVE-2026-39987 Exploited Within 10 Hours of Disclosure -- The Hacker News (2026) -- https://thehackernews.com/2026/04/marimo-rce-flaw-cve-2026-39987.html

[43] NetScaler ADC and NetScaler Gateway Security Bulletin for CVE-2026-3055 and CVE-2026-4368 -- Citrix (2026) -- https://support.citrix.com/external/article/CTX696300/netscaler-adc-and-netscaler-gateway-secu.html

[44] CVE-2026-3055: Citrix NetScaler ADC and NetScaler Gateway Out-of-Bounds Read -- Rapid7 (2026) -- https://www.rapid7.com/blog/post/etr-cve-2026-3055-citrix-netscaler-adc-and-netscaler-gateway-out-of-bounds-read/

[45] CVE-2026-3055: Citrix NetScaler Memory Overread -- Horizon3.ai (2026) -- https://horizon3.ai/attack-research/vulnerabilities/cve-2026-3055/

[46] CVE-2026-3055: Citrix NetScaler Memory Leak Actively Exploited -- CrowdSec (2026) -- https://www.crowdsec.net/vulntracking-report/cve-2026-3055-citrix-netscaler-memory-leak-exploitation

[47] Rapid Reaction - Citrix NetScaler ADC and NetScaler Gateway Memory Overread Vulnerability (CVE-2026-3055) -- watchTowr (2026) -- https://watchtowr.com/resources/rapid-reaction-to-cve-2026-3055/

[48] CVE-2026-33824: Windows IKE Extension RCE Vulnerability -- SentinelOne Vulnerability Database (2026) -- https://www.sentinelone.com/vulnerability-database/cve-2026-33824/

[49] Security Update Guide -- CVE-2026-33824 -- Microsoft Security Response Center (2026) -- https://msrc.microsoft.com/update-guide/vulnerability/CVE-2026-33824

[50] CVE-2026-9198: Langflow RCE Vulnerability -- SentinelOne Vulnerability Database (2026) -- https://www.sentinelone.com/vulnerability-database/cve-2026-9198/

[51] CVE-2026-18556 -- N-able N-central: Authentication bypass -- Rapid7 Vulnerability Database (2026) -- https://www.rapid7.com/db/vulnerabilities/cve-2026-18556/

[52] CISA Adds Exploited N-able N-central Flaw to KEV After Customer Compromises -- The Hacker News (2026) -- https://thehackernews.com/2026/08/cisa-adds-exploited-n-able-n-central.html