Metasploit Weaponizes Check Point SmartConsole, Tenable Security Center, and the Langflow/Flowise AI-Agent Frameworks

A Defender's Technical Briefing

Scope note: This report covers exploit modules Rapid7 added to the open-source Metasploit Framework for four security-tooling and AI-orchestration products between April 2025 and August 28, 2026: Check Point SmartConsole, Tenable Security Center, Langflow, and Flowise. All four are now targets with public, weaponized exploit code in the world's most widely used penetration-testing framework — meaning any attacker with `msfconsole` and network access can reproduce these compromises without writing exploit code themselves. This report explains the underlying weaknesses, the exact affected/fixed version matrices, what is verifiably known about in-the-wild exploitation, and what defenders can do today. It does not include exploit code or step-by-step attack instructions.

DEFENDER'S TECHNICAL BRIEFING

Metasploit module data, CVE identifiers, and KEV entries reflect public information as of August 30, 2026.

Contents

1. Why this matters: the pattern behind these four modules

All four products share a structural property that makes them attractive Metasploit targets: they are **trusted, privileged infrastructure** that organizations expose to reach broad swaths of their environment. A security-management console has standing credentials and control-plane access to every managed firewall. A vulnerability-management platform aggregates authenticated scan data (and often scan *credentials*) for an entire estate. An AI-agent orchestration platform is, by design, a code-execution engine wired to LLM API keys, databases, and downstream tools. Compromising any of them yields a force-multiplying foothold rather than a single host.

Figure A. Four unrelated products, one shared risk shape: each is trusted management/orchestration infrastructure whose compromise yields broad downstream access.
Original illustration, derived from Section 1 of this report.

Rapid7's own August 28, 2026 “Metasploit Wrap-Up” post bundled new modules for Check Point SmartConsole, two Tenable Security Center vulnerabilities, and a Flowise vulnerability in the same release — a coincidence of scheduling, not a shared vulnerability class, but a useful signal that Metasploit's module authors (Rapid7 staff and outside contributors alike) are actively targeting

management/orchestration-plane software as a category (Rapid7, "Metasploit Wrap-Up," Aug 28, 2026). Separately, Langflow and Flowise — both open-source, low-code frameworks for building LLM-powered "agentic" pipelines — have each accumulated a *cluster* of critical CVEs and Metasploit modules over the past 18 months, illustrating a distinct and recurring root cause: these frameworks treat arbitrary code execution as a first-class *feature* (running user- or LLM-generated Python/JavaScript to build dynamic workflows), and the security boundary meant to contain that feature has repeatedly failed.

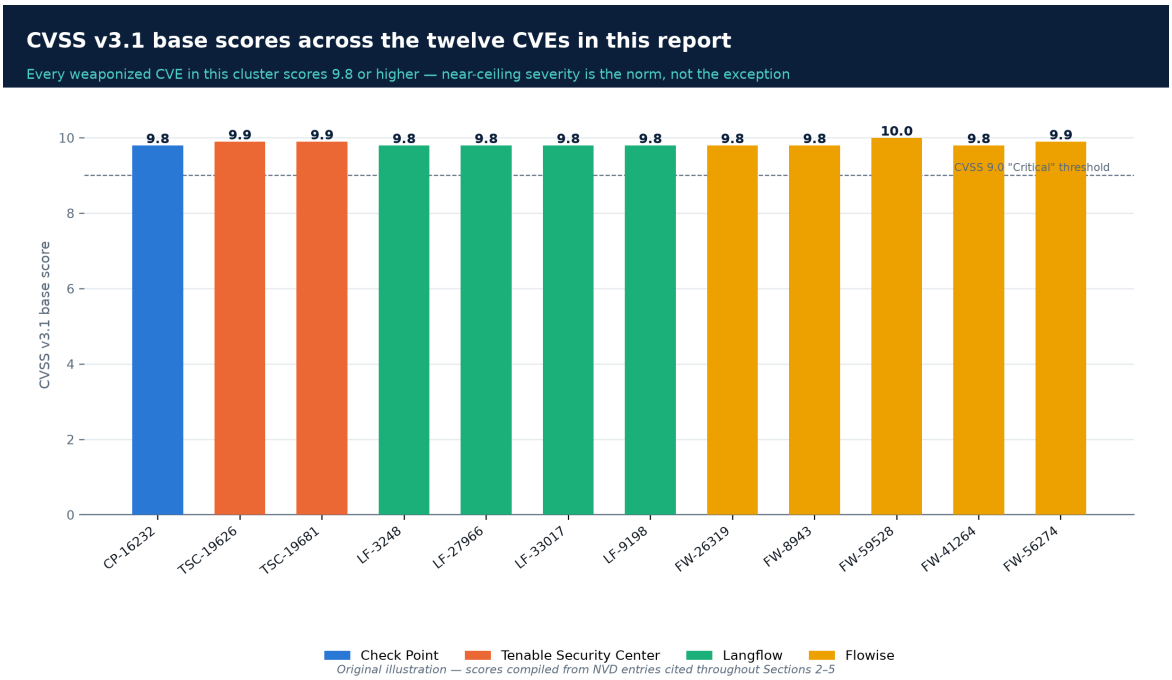

Figure B. Severity is not the differentiator in this cluster — every weaponized CVE discussed in this report scores 9.8 or above.

Original illustration, scores compiled from the NVD entries cited in Sections 2–5.

The rest of this report is organized product-by-product, followed by consolidated detection guidance, interim mitigations, and open questions.

2. Check Point SmartConsole — CVE-2026-16232

2.1 What the vulnerability is

CVE-2026-16232 is an **authentication-bypass vulnerability** (CWE-287, Improper Authentication) in the login process used by Check Point SmartConsole to authenticate against Check Point Security Management Servers and Multi-Domain Security Management Servers (MDS) — i.e., the centralized consoles that administer Check Point firewalls. NVD scores it CVSS v3.1 **9.8 (Critical)** (`CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H`) and CVSS v4.0 **9.3 (Critical)** — NVD, CVE-2026-16232. Note that Check Point's own advisory and some secondary write-ups cite a slightly different consolidated score of 9.1; defenders should treat this as the same critical-severity finding described with different scoring methodologies rather than two different bugs.

Root cause. Per Rapid7's independent technical analysis, the flaw lives in the trust boundary between a SmartConsole GUI client and the management server. The server accepts an attacker-supplied Secure Internal Communication (SIC) distinguished name as a proxy for identity, instead of cryptographically binding the caller's identity to an authenticated peer certificate. An unauthenticated network attacker can therefore obtain an **application login token**, use that token to mint a valid single-sign-on ticket, and redeem the ticket through the management server's CPM SOAP service to log in to SmartConsole with **full administrative privileges** — no credentials, no prior access required (Rapid7, "Check Point SmartConsole Authentication Bypass — Technical Analysis," CVE-2026-16232; Check Point sk185169). Full administrative access to a Security Management Server means an attacker can rewrite firewall policy across every gateway the server manages — a direct path to disabling network defenses at scale.

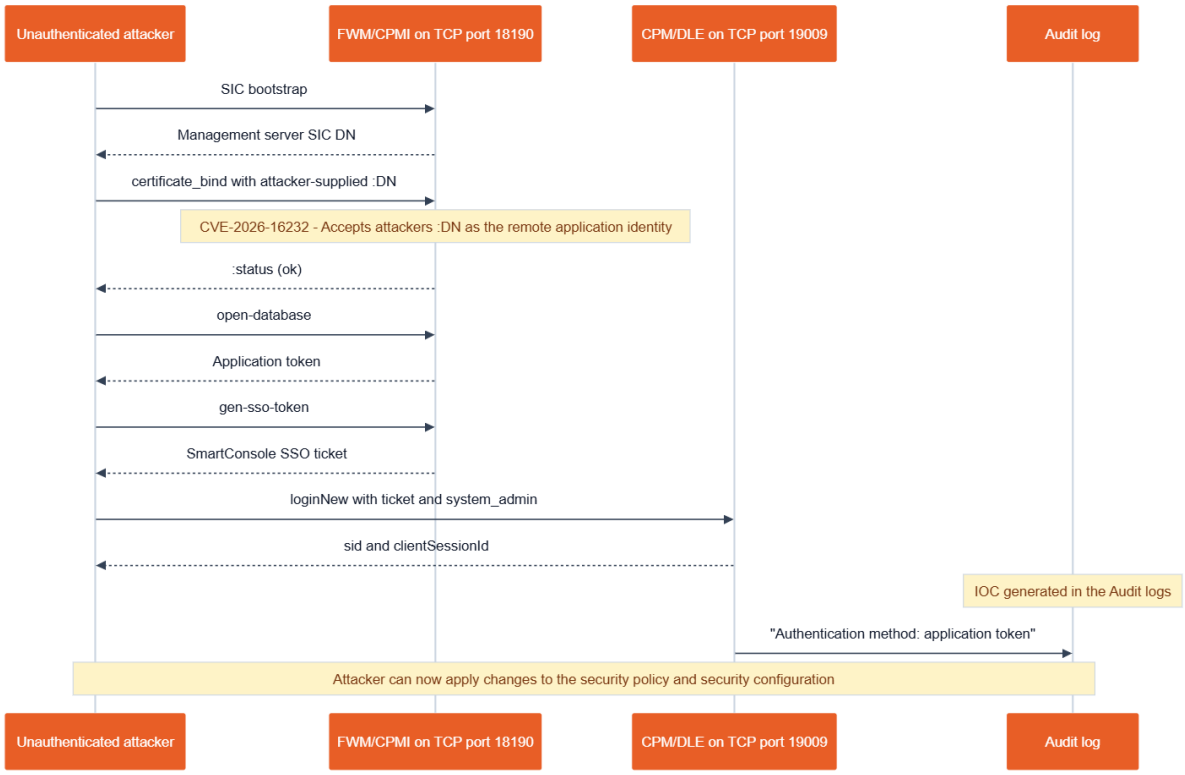

Figure 1. The CVE-2026-16232 exploitation flow: an unauthenticated attacker bootstraps SIC, supplies a forged distinguished name, and rides it through application-token issuance and SSO-ticket redemption to a full SmartConsole login. The final step leaves the "Authentication method: application token" IOC in the audit log.

Credit: Rapid7, "Check Point SmartConsole Authentication Bypass — Technical Analysis (CVE-2026-16232)."

Attack prerequisites. Two conditions must hold for remote exploitation: (1) the Security Management/MDS server's management interface must be reachable from the attacker's network position — in the worst case, directly exposed to the internet — and (2) the "Trusted Clients" (GUI-client access control) setting must be left at its permissive default of "Any" rather than scoped to specific administrator IP ranges. Check Point and Rapid7 both note the second condition is a commonly encountered, not a rare, misconfiguration.

2.2 Affected and fixed version matrix

Per Check Point's advisory sk185169, the affected product line spans Security Management Server / Multi-Domain Server releases **R77.30, R80, R80.10, R80.20, R80.30, R81, R81.10, R81.20, R82, and R82.10**. Confirmed fixed builds (Jumbo Hotfix Accumulator, "JHA," "Take" numbers) are published for the actively supported tracks only:

Track	Status	Fixed at
R82.10	Vulnerable	Jumbo Hotfix Accumulator Take 36 and later
R82	Vulnerable	Jumbo Hotfix Accumulator Take 118 and later
R81.20	Vulnerable (confirmed exploited pre-patch)	Jumbo Hotfix Accumulator Take 158 and later
R81.10, R81, R80.40/30/20/10/00, R77.30	Listed as affected in the advisory's version list	No fixed-build number published for these end-of-support tracks — treat as “affected, unsupported”

Source: Check Point sk185169; corroborated by Rapid7's technical analysis and Emergent Threat Response report. Organizations on end-of-support branches have no vendor-supplied fix path other than upgrading to a supported major version.

2.3 Is it being exploited in the wild?

Yes — confirmed by the vendor itself, independently corroborated by CISA. Check Point's own security advisory states the company became aware of active exploitation, describing it as affecting **“a handful of customers with specific configurations”** in which the management server was **“exposed directly to the internet without IP restrictions”** (Check Point blog, "Security Advisory: Action Required"). Check Point states Smart-1 Cloud customers were already protected and that all affected customers were individually notified. CISA independently corroborated active exploitation by adding CVE-2026-16232 to its **Known Exploited Vulnerabilities (KEV) catalog on July 22, 2026**, with a compressed remediation deadline of **July 25, 2026** — a three-day window that federal agencies must meet under CISA Binding Operational Directive 26-04, which replaced the older BOD 22-01 and reserves the shortest (3-day) remediation tier for vulnerabilities that combine KEV status, public exposure, automatable exploitation, and total-control impact.

This is a case where the disclosure order was: **zero-day exploitation in the wild → vendor patch and advisory (July 22, 2026, same day) → CISA KEV addition (same day) → independent researcher PoC (Rapid7's Stephen Fewer, ~July 23–28, 2026) → full weaponized Metasploit module (same researcher, merged and announced August 28, 2026)** — roughly five weeks from confirmed zero-day use to a productized, one-command exploit module in the world's most widely deployed pentest framework. Coverage of the exploited zero-day was corroborated by The Hacker News, BleepingComputer, and Cybersecurity Dive.

2.4 Metasploit module

- **Module path:** `exploit/linux/misc/checkpoint_smartconsole_cve_2026_16232_rce`
- **Author:** Stephen Fewer (Rapid7, handle `sfewer-r7`)
- **Merged via:** rapid7/metasploit-framework PR #21731
- **Announced:** Rapid7 Metasploit Wrap-Up, August 28, 2026, described there as exploiting “an authentication bypass in the SmartConsole login process affecting Check Point Security Management Server and Multi-Domain Security Management Server (MDS).”

2.5 Detection guidance

Check Point's advisory and Rapid7's Emergent Threat Response (ETR) report jointly publish the following independently verifiable artifacts:

- **Audit-log signature:** search SmartConsole/Security Management audit logs for the string `Authentication method: application token`. Legitimate administrator logins authenticate with credentials (username/password, certificate, or SSO), not a bare application token — its presence in an audit-log entry is a strong indicator this technique was used.
- **Historical attacker source IPs** (Check Point-published, relayed via Rapid7's ETR — treat as a non-exhaustive, campaign-specific list, not a complete blocklist): `151.241.99[.]207`, `151.241.99[.]233`, `158.62.198[.]182`, `192.142.10[.]99`, `139.28.37[.]250`, `194.213.18[.]137`. Absence of these specific IPs in your logs does **not** confirm you were not targeted.
- **Configuration audit:** verify `Manage & Settings > Permissions & Administrators > Trusted Clients` is scoped to specific administrator IP addresses/subnets rather than “Any” — this doubles as both the exploitation precondition and a detective control.
- **Vulnerability scanning:** Greenbone added an authenticated/unauthenticated remote version-detection check to its OpenVAS Enterprise Feed for vulnerable Gaia OS builds (Greenbone).
- **Validation tooling:** Rapid7's published Python PoC can be used by defenders (in an authorized testing context) to confirm whether a given management server is patched — it is dual-use with the same code path the Metasploit module exploits, so treat it with the same operational caution as any exploit code.

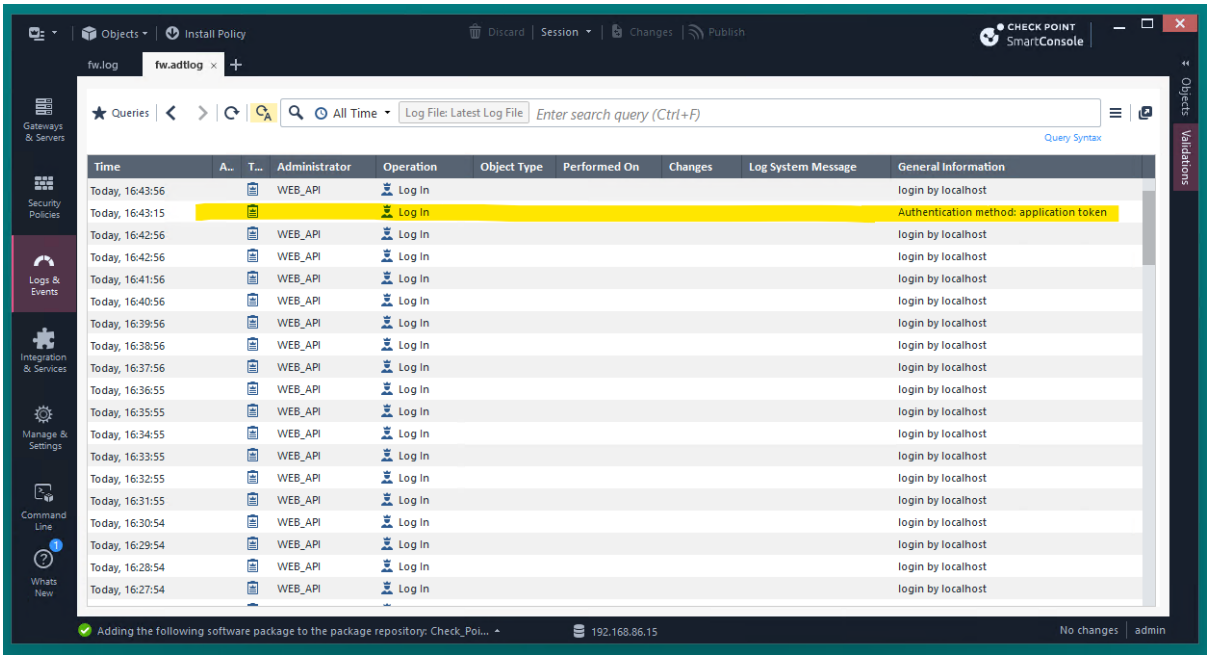

Figure 2. The audit-log artifact left behind by exploitation: a login entry whose General Information field reads “Authentication method: application token” — the primary detection signature for CVE-2026-16232.

Credit: Rapid7, "Check Point SmartConsole Authentication Bypass — Technical Analysis (CVE-2026-16232)."

2.6 Interim mitigations for teams that cannot patch immediately

Per Check Point's own advisory:

1. **Restrict Trusted Clients** to the specific IP addresses/subnets used by legitimate SmartConsole administrators — do not leave this set to “Any.”
2. **Firewall the management interface.** Ensure the Security Management/MDS server's management IP is not reachable from the general internet; restrict inbound access to trusted administrator networks only (VPN or jump-host access is preferable to direct exposure).
3. **Verify implied rules for control connections** are enabled and correctly scoped, per Check Point's guidance.
4. Prioritize patching to the Jumbo Hotfix Accumulator Take listed in §2.2 for your track; for end-of-support tracks with no published fix, plan an accelerated upgrade to a supported release.

3. Tenable Security Center — CVE-2026-19626 and CVE-2026-19681

3.1 What the vulnerabilities are

Tenable's August 13, 2026 advisory **TNS-2026-22** ("[R1] Security Center Version 6.9.0 Fixes Multiple Vulnerabilities") bundles fixes for a batch of first-party Security Center flaws plus updated third-party components (curl, underscoreJS). Two of the first-party CVEs from this batch were separately weaponized as Metasploit modules.

CVE-2026-19626 — "Report Charting RCE"

CWE-95, Improper Neutralization of Directives in Dynamically Evaluated Code (eval injection). CVSS v3.1 **9.9 (Critical)**, CVSS v4.0 **9.4 (Critical)**. Per Tenable's and NVD's descriptions, "an authenticated, non-administrative user could exploit this issue by supplying specially crafted input that is later processed unsafely during server-side report rendering, resulting in arbitrary code execution with the privileges of the service account." In practice, a low-privileged Security Center user who can create or customize a report (e.g., a pie-chart/group-name label) can inject code that the server evaluates when the report is rendered (NVD, CVE-2026-19626; Tenable TNS-2026-22).

CVE-2026-19681 — "SCAP Audit File Command Injection"

CWE-78, OS Command Injection. CVSS v3.1 **9.9 (Critical)**, CVSS v4.0 **9.4 (Critical)** (same vector shape as above). Per NVD: "An authenticated command injection vulnerability exists in Security Center related to file upload processing. An attacker could exploit this issue by uploading a specially crafted file, potentially resulting in arbitrary command execution on the underlying operating system" — consistent with third-party reporting describing this as exploitable via a crafted SCAP audit-file upload (NVD, CVE-2026-19681).

Both require **authentication** (low-privileged, non-administrative access is sufficient) — these are not unauthenticated, internet-scannable bugs, which meaningfully changes the threat model relative to the Check Point and Langflow/Flowise flaws discussed elsewhere in this report. The practical risk vector is an insider, a compromised low-privilege credential (e.g., phished analyst account), or a pivot from an adjacent compromised system that already holds Security Center credentials.

3.2 Affected and fixed version matrix

Both CVEs affect **Tenable Security Center 6.8.0 and earlier**; both are fixed in **Security Center 6.9.0**. Tenable's advisory offers no interim workaround — the stated remediation is the full upgrade to 6.9.0. This is a distinct, later advisory from Tenable's earlier July 20, 2026 advisory **TNS-2026-19**, which shipped a standalone patch (SC202607.1) for a different set of CVEs affecting Security Center 6.6.0, 6.7.2, and 6.8.0 — defenders should not conflate the two advisories when tracking patch status (Canadian Centre for Cyber Security advisory AV26-724).

3.3 Is it being exploited in the wild?

No confirmed evidence found as of this writing (August 30, 2026). A direct query of CISA's KEV JSON feed for CVE-2026-19626 and CVE-2026-19681 returned no matches. No GreyNoise, Shadowserver, or vendor threat-intelligence report describing active exploitation was located. However, defenders should treat this as an elevated near-term risk rather than a closed issue: a public technical write-up discussing exploitation of the eval-injection flaw has already circulated on a community security forum (KSEC forum thread), and the Metasploit modules themselves became public roughly 12 days after the CVEs were disclosed — a well-documented pattern (see the Langflow and Flowise sections below) in which public weaponization sharply narrows the time between disclosure and mass exploitation attempts.

3.4 Metasploit modules

CVE	Module	Author	PR
CVE-2026-19626	exploit/linux/http/tenable_sc_report_charting_rce_cve_2026_19626	h00die	#21820
CVE-2026-19681	exploit/linux/http/tenable_sc_auditfile_cmdinject_cve_2026_19681	h00die	#21802

Both announced in Rapid7's Metasploit Wrap-Up, August 28, 2026, which describes each as “an authenticated remote code execution vulnerability” against Tenable Security Center.

3.5 Detection guidance and interim mitigations

Tenable has not published dedicated log signatures or IOCs for these two CVEs. In their absence, the following controls follow directly from the vulnerability mechanics and are consistent with standard eval-injection/command-injection hardening (labeled here as inference, not vendor-sourced guidance, since no such guidance has been published):

- **Restrict report-authoring and SCAP audit-file upload privileges** to trusted administrators only; do not grant “create/customize report” or “upload audit file” rights broadly to non-administrative analyst roles until patched.
- **Monitor Security Center's application/service-account process tree** for anomalous child-process spawning that coincides with report-generation jobs or SCAP audit-file uploads — a service account spawning a shell or unexpected binary during what should be a report-render or file-parse operation is the expected artifact of successful exploitation of either flaw.
- **Network-segment the Security Center management interface** the same way you would any other privileged vulnerability-management console — it holds scan credentials for the rest of your estate.
- **Patch to 6.9.0.** No standalone hotfix exists for this batch; the full version upgrade is the only vendor-sanctioned remediation (Tenable TNS-2026-22).

4. Langflow — a cluster of critical CVEs and the first documented “agentic ransomware”

Langflow is an open-source, drag-and-drop framework (now distributed as “IBM Langflow OSS” in some advisories) for building LLM-powered agent workflows. It has accumulated at least seven CVEs with public exploitation evidence or Metasploit weaponization since April 2025, six of which CISA has added to its KEV catalog — an unusually dense cluster for a single product in this timeframe.

4.1 CVE-2025-3248 — the original unauthenticated RCE, and its exploitation history

Root cause. Langflow's `/api/v1/validate/code` endpoint accepts attacker-supplied Python, uses the `ast` module to statically filter out obviously dangerous constructs (imports, function definitions), and then passes the “validated” code through `compile()`/`exec()`. Horizon3.ai's disclosure explains the filter's fatal flaw: **Python decorators and default-argument expressions are evaluated immediately at function-definition time**, not at call time — so a payload wrapped in a decorator (or hidden in a default argument) executes *during the validation step itself*, before any call-time filtering could apply (Horizon3.ai, "Unsafe at Any Speed"). Horizon3 reported the flaw existed unnoticed in Langflow's codebase for roughly two years before disclosure.

CVE and versions. CVE-2025-3248, CWE-94 (Code Injection) and CWE-306 (Missing Authentication), CVSS v3.1 **9.8 (Critical)** per NVD. Affects **all Langflow versions before 1.3.0**; fixed in **1.3.0** (released March 31, 2025). Disclosure timeline per Horizon3: reported to the vendor February 22, 2025; patch merged March 4–5, 2025; 1.3.0 shipped March 31, 2025; CVE assigned April 3, 2025; Horizon3's public write-up and exploit followed April 9, 2025.

Metasploit module. `exploit/multi/http/langflow_unauth_rce_cve_2025_3248`, authored by **Naveen Sunkavally** (Horizon3.ai, vulnerability research) and **Takahiro Yokoyama** (module implementation), disclosure date April 9, 2025, ranked “Excellent” (crash-safe, repeatable sessions) — Rapid7 module database entry.

CISA KEV. Added **May 5, 2025**, federal remediation due May 26, 2025; corroborated by The Hacker News.

In-the-wild exploitation — three independently documented campaigns

This is the best-evidenced in-the-wild case in this report:

1. **Flodrix botnet.** Trend Research documented threat actors scanning for internet-exposed Langflow instances and using CVE-2025-3248 to deploy **Flodrix**, a DDoS-capable botnet in the LeetHozer/Mirai lineage (Trend Micro Research).
2. **SANS Internet Storm Center** documented exploitation attempts against CVE-2025-3248 within days of Horizon3's April 9, 2025 public disclosure, indicating rapid mass scanning followed public weaponization almost immediately (SANS ISC diary, April 12, 2025).

3. **JADEPUFFER — the first documented fully agentic ransomware attack.** Sysdig's Threat Research Team published a detailed account of an intrusion in which CVE-2025-3248's unauthenticated code-execution flaw in Langflow's code-validation endpoint served as the confirmed initial access vector. What followed was executed by an autonomous LLM-driven agent, not a human operator: host and process reconnaissance, credential harvesting across cloud providers/API keys/databases/crypto wallets, MinIO object-storage enumeration via default credentials, lateral movement to a production database server, a Nacos configuration-service takeover via CVE-2021-29441 combined with JWT forgery, installation of a self-correcting database backdoor, AES encryption of 1,342 configuration items, destructive database wiping with an inserted ransom note, and crontab-based persistence. Sysdig captured over 600 distinct payloads containing natural-language reasoning and self-narrating annotations from the agent, including a 31-second autonomous recovery from a failed login attempt and pre-checks for container escape via `docker.sock` probing (Sysdig, "JADEPUFFER"). This finding was independently covered by The Hacker News, BleepingComputer, SecurityWeek, and CyberScoop. For defenders, the significance is less about Langflow specifically and more about a demonstrated capability: a still-unpatched, 15-month-old CVE remained a viable initial-access vector for a fully autonomous post-exploitation chain in mid-2026.

Scanning scale. Multiple secondary sources estimate several thousand internet-exposed Langflow instances discoverable via Shodan/Censys/FOFA-style reconnaissance (figures in the low thousands to tens of thousands vary by source and survey date); treat exact counts as approximate and time-bound rather than a fixed figure (saassentinel.com survey).

4.2 Subsequent Langflow CVEs and KEV entries (2026)

Langflow's CVE cadence accelerated through 2026. The table below lists every additional Langflow CVE independently confirmed for this report, in each case cross-checked against NVD and/or the CISA KEV JSON feed directly:

CVE	Class / root cause	CVSS	Affected → Fixed	CISA KEV	Metasploit module
CVE-2026-27966	CWE-94. CSV Agent component hardcoded `allow_dangerous_code=True`, exposing LangChain's `python_repl_ast` tool to unauthenticated prompt injection.	9.8	< 1.8.0 → 1.8.0 (adds explicit `BoolInput` gate)	Not found in KEV as of Aug 30, 2026	langflow_rce_cve_2026_27966, author Takahiro-Yoko, PR #21260
CVE-2026-33017	CWE-94. `POST /api/v1/build_public_tmp/{flow_id}/flow` accepts attacker-controlled flow data instead of the stored, trusted flow — arbitrary Python passed straight to `exec()` with zero sandboxing.	9.8	≤ 1.8.2 (and dev builds through 1.9.0.dev11) → 1.9.0	Yes — added March 25, 2026, due April 8, 2026	langflow_unauth_rce_cve_2026_33017, author Richard Howe (rhowe425), disclosure March 20, 2026, PR #21700
CVE-2025-34291	CWE-346, Origin Validation Error. Chains a misconfigured CORS policy, a refresh-token cookie scoped for cross-site delivery, and missing CSRF protection to achieve full account takeover and RCE from a single victim browser visit — no stolen credentials required.	CVSS v4.0 9.4	≤ 1.6.9	Yes — added May 21, 2026, due June 4, 2026; active exploitation reported by CrowdSec researchers beginning January 23, 2026, months before the KEV listing	Not confirmed
CVE-2026-55255	Authorization bypass through user-controlled key (IDOR) — an authenticated attacker can execute another user's flow by specifying the victim's flow ID.	Not independently confirmed via NVD in this research pass	Not confirmed in this research pass	Yes — added July 7, 2026, due July 10, 2026	Not confirmed
CVE-2026-0770	“Inclusion of Functionality from Untrusted Control Sphere” — remote attackers can execute arbitrary code.	Not independently confirmed via NVD in this research pass	Fix referenced at the Langflow v1.9.0 release	Yes — added July 21, 2026, due July 24, 2026	Not confirmed
CVE-2026-9198	Chained authentication-bypass + code-injection: `GET /api/v1/auto_login` mints a SUPERUSER JWT to any unauthenticated caller whenever `AUTO_LOGIN` is enabled (Langflow's default), which is then used to reach `/api/v1/validate/code` and execute arbitrary Python via `exec()` with no sandbox.	9.8	1.0.0–1.10.0 → 1.10.1	Yes — added August 4, 2026, due August 7, 2026; listed under vendor “IBM”	langflow_unauth_rce_cve_2026_9198, author Richard Howe (rhowe425), disclosure July 17, 2026

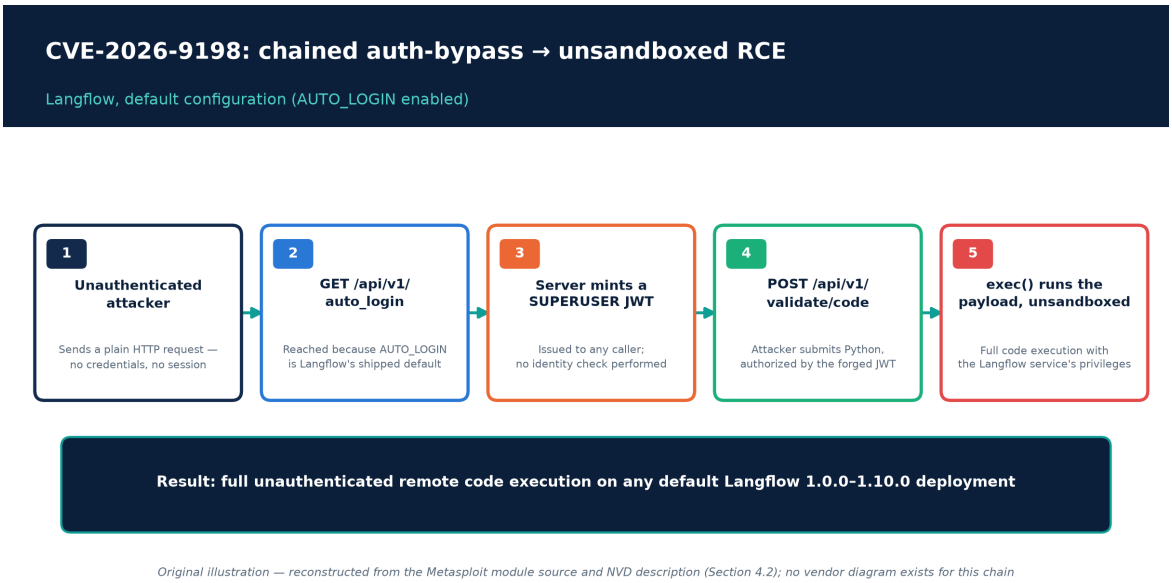

Figure C. How CVE-2026-9198 chains a default-on convenience feature into unauthenticated remote code execution. Original illustration, reconstructed from the Metasploit module source and NVD description (Section 4.2); no vendor diagram exists for this chain.

A note on incomplete patches. JFrog Security Research separately reported that the version Langflow shipped as the "fix" for CVE-2026-33017 remained exploitable via a variant of the same code path — a caution against treating a single point release as closing the issue without independent verification (JFrog Security Research).

Detection guidance across the Langflow CVE set

- Monitor/alert on unauthenticated or unexpected `POST` requests to `/api/v1/validate/code`, `/api/v1/build_public_tmp/{flow_id}/flow`, and `GET` requests to `/api/v1/auto_login` from source addresses that are not your known administrative ranges.
- A Nuclei detection template exists (referenced in Horizon3's original disclosure) that probes `/api/v1/validate/code` with decorator-based payloads, useful for authorized asset-inventory scanning of CVE-2025-3248 exposure.
- If Flodrix/Gafgyt-style botnet activity is a concern, monitor for outbound connections consistent with Mirai-family C2 beaconing originating from Langflow hosts, and treat any previously internet-exposed pre-1.3.0 instance as a candidate for compromise assessment, not just patching.
- Given the `AUTO_LOGIN`-default root cause of CVE-2026-9198, explicitly audit whether `AUTO_LOGIN` is enabled in your deployment and disable it unless strictly required, independent of patch status.

5. Flowise — a comparable AI-agent-framework cluster, now an end-of-life product

Flowise is a similar open-source, low-code AI-agent-building platform. Independently of Langflow, it has accumulated its own critical-CVE cluster, several of which were confirmed under active, concurrent exploitation by external researchers — and the project itself was **archived and formally sunset by its maintainers in August 2026**, a fact with direct consequences for long-term patch availability.

5.1 The CVE set

CVE	Class / root cause	CVSS	Affected → Fixed	Metasploit module
CVE-2025-26319	Pre-authentication arbitrary file upload via `/api/v1/attachments` — no restriction on file extension, MIME type, or content.	9.8	≤ 2.2.6 → 2.2.7-patch.1	Not confirmed
CVE-2025-8943	CWE-306/CWE-862. Custom MCP feature spawns OS processes (e.g., via `npx`) through `StdioClientTransport` without authentication enforcement; pre-3.0.1, default installs run with no auth unless explicitly configured.	9.8	< 3.0.1 → 3.0.1	flowise_custommcp_rce, authors Assaf Levkovich and Valentin Lobstein
CVE-2025-59528	CWE-94. `convertToValidJSONString()` in the CustomMCP node passes attacker-controlled configuration straight into JavaScript's `Function()` constructor (functionally `eval`) at `POST /api/v1/node-load-method/customMCP`.	10.0 (maximum)	2.2.7-patch.1–3.0.5 → 3.0.6 (replaces `Function()` with `JSON5.parse()`)	flowise_js_rce, authors Kim SooHyun, nltt0, Valentin Lobstein
CVE-2026-41264	Sandbox-escape RCE via CSV Agent prompt injection — `CSV_Agents.run()` evaluates LLM-generated Python with no sandboxing. Discovered by Trend Micro's Zero Day Initiative.	9.8	1.3.0–3.0.13 → 3.1.0	flowise_auth_rce_cve_2026_41264, authors zdi-disclosures and Takahiro Yokoyama, disclosure April 22, 2026
CVE-2026-56274	CWE-78. Multiple OS command-injection paths in the Custom MCP Server feature: incomplete command-flag validation permits abuse of the `npx --yes` flag, plus a regex bypass in local-file-access restrictions.	v3.1: 9.9; v4.0: 8.7 (both legitimate NVD scores under different CVSS versions)	2.2.7–3.1.1 → 3.1.2	flowise_mcp_rce, authors cn-panda and ABDUL JAFAROV, disclosure June 23, 2026, PR #21616

Both `flowise_js_rce` and `flowise_custommcp_rce` were added to Metasploit in the same PR #20705 (opened November 18, 2025; merged November 21, 2025; author Chocapikk), which also introduced a shared HTTP mixin library for Flowise targeting. `flowise_mcp_rce` (CVE-2026-56274) was announced in the same August 28, 2026 Metasploit Wrap-Up that carried the Check Point and Tenable modules discussed above.

A module authentication detail relevant to detection: for Flowise versions before 3.0.1, these exploits can run **fully unauthenticated** if the operator never configured `FLOWISE_USERNAME`/`FLOWISE_PASSWORD`; from 3.0.1 onward, JWT enforcement means the modules require valid `FLOWISE_EMAIL`/`FLOWISE_PASSWORD` credentials — reinforcing that "is auth enabled" is not a reliable indicator of safety for pre-3.0.1 deployments left at defaults.

5.2 Is it being exploited in the wild?

Yes, for at least three of the above CVEs, concurrently. VulnCheck (researcher Caitlin Condon) detected active exploitation of **CVE-2025-59528** beginning **April 7, 2026**, originating from a single Starlink-attributed IP address conducting low-volume but sustained, automated scanning; the same reporting states **CVE-2025-8943** and **CVE-2025-26319** were being exploited concurrently, and cites over **12,000 internet-exposed Flowise instances** at the time of the report (The Hacker News, "Flowise AI Agent Builder Under Active Attack," April 7, 2026). A Cloud Security Alliance research note published the same week corroborates the findings and specifically recommends that any exposed pre-3.0.6 instance be treated as **already compromised** — i.e., that credential rotation, not just patching, is the appropriate remediation step (Cloud Security Alliance research note). Notably, CVE-2025-59528 had been publicly disclosed roughly six months before this active-exploitation wave began, illustrating that "old" AI-framework CVEs remain live threats as exposed instance counts stay high.

None of the Flowise CVEs above appear in CISA's KEV catalog as of August 30, 2026 (confirmed by a direct query of the CISA KEV JSON feed for all six Flowise CVE IDs — no matches). This does not indicate the exploitation described above is unconfirmed; KEV listing reflects CISA's own confirmation process and federal-agency relevance criteria, and the VulnCheck/CSA reporting is independent, primary evidence in its own right.

5.3 Detection guidance

- **Network signature:** monitor `POST` requests to `/api/v1/node-load-method/customMCP` for JavaScript payload indicators such as `process.mainModule`, `child_process`, `require`, or `execSync`/`exec` in the request body — this is the CVE-2025-59528 injection point.
- **Commercial IPS coverage:** SonicWall has published IPS signatures **21519** and **21918** ("FlowiseAI Flowise CustomMCP Remote Code Execution 1/2") specifically for CVE-2025-59528.
- **Nuclei template:** a community-maintained template exists for the earlier CVE-2024-31621 case-sensitivity authentication bypass, useful as a baseline exposure/asset-discovery scan even though that specific CVE predates the modules covered in this report.
- Given the CSA guidance above, **any organization that discovers a pre-3.0.6 Flowise instance that was internet-exposed during the April 2026 exploitation window should conduct a compromise assessment and rotate every credential the instance held** (LLM provider API keys, database credentials, cloud service-account keys), not merely apply the patch.

5.4 Critical operational fact: Flowise is now an archived, unsupported project

Independently confirmed for this report: FlowiseAI announced a shutdown timeline in which **active feature development and pull-request review ceased on July 29, 2026** (“code freeze”), the GitHub repository transitioned to **public-archive (read-only) status around August 10–13, 2026**, and **core-team community support on Discord/GitHub concluded on August 31, 2026** (FlowiseAI sunset announcement). The announcement does not commit to any further official security patching. The Apache-2.0-licensed source remains available for self-hosters to fork, and the FlowiseAI team has explicitly encouraged community forks to take over maintenance — but as of this writing, **no organization should assume future Flowise CVEs will receive an official vendor patch**. For any team still running Flowise in production, this elevates the priority of migrating off the platform (or adopting/vetting a maintained community fork) above simply keeping current with the last official release (3.1.2).

6. Cross-cutting analysis: what these four cases teach about AI-agent-framework risk specifically

The Check Point and Tenable cases are classic management-plane vulnerabilities (auth-bypass, eval/command injection in a privileged service) of a kind security teams have handled for decades. Langflow and Flowise represent a newer and structurally different failure mode worth calling out explicitly for defenders evaluating AI-agent platforms generally (LangGraph, LangChain-based custom deployments, and comparable low-code AI orchestration tools share the same underlying risk shape even where no CVE yet exists):

1. **Arbitrary code execution is the product, not a bug.** Langflow's `validate/code` endpoint and Flowise's CustomMCP/CSV Agent nodes exist specifically so that users (or the LLM itself, via agentic tool use) can run dynamically generated code as part of a workflow. Every mitigation in this class of product is therefore a *sandbox* or *authorization gate* around a feature whose entire purpose is to execute code — a much harder security property to guarantee than “this endpoint should never execute code at all.”
2. **Filters built on static analysis of dynamic languages are fragile.** Both the original Langflow flaw (decorators/default arguments evaluated at “validation” time) and the Flowise flaw (JavaScript `Function()` constructor abuse) exploited the gap between what a security filter can see statically and what the language actually executes. This is a recurring, language-level failure mode, not an implementation bug specific to either project.
3. **Default-enabled dangerous features compound the problem.** CVE-2026-27966's root cause was a hardcoded `allow_dangerous_code=True`; CVE-2026-9198 depended on `AUTO_LOGIN` being enabled by default; CVE-2025-8943 depended on authentication being off by default pre-3.0.1. In each case, the vulnerable configuration was the out-of-the-box state, not an unusual hardening lapse — meaning exposure scales with the number of default deployments, not the number of misconfigured ones.
4. **Prompt injection is now a code-execution vector, not just a content-manipulation nuisance.** CVE-2026-27966 and CVE-2026-41264 both use prompt injection against an LLM agent (CSV Agent) as the delivery mechanism for what is, functionally, a remote code execution exploit. Traditional web-application threat models that treat “user input” and “attacker input” as the same category do not cleanly map onto agentic systems where an LLM's *output*, shaped by untrusted *input data* (a CSV file, a document, a webpage the agent reads), becomes executable.
5. **The disclosure-to-exploitation and disclosure-to-weaponization windows are compressing.** CVE-2026-33017 saw active exploitation within roughly 20 hours of disclosure, before any public PoC existed. CISA's own BOD 26-04 cites this exact dynamic — the collapse of the disclosure-to-weaponization window from months to hours — as its rationale for replacing static CVSS-based remediation deadlines with a risk-tiered model that can mandate 3-day remediation.

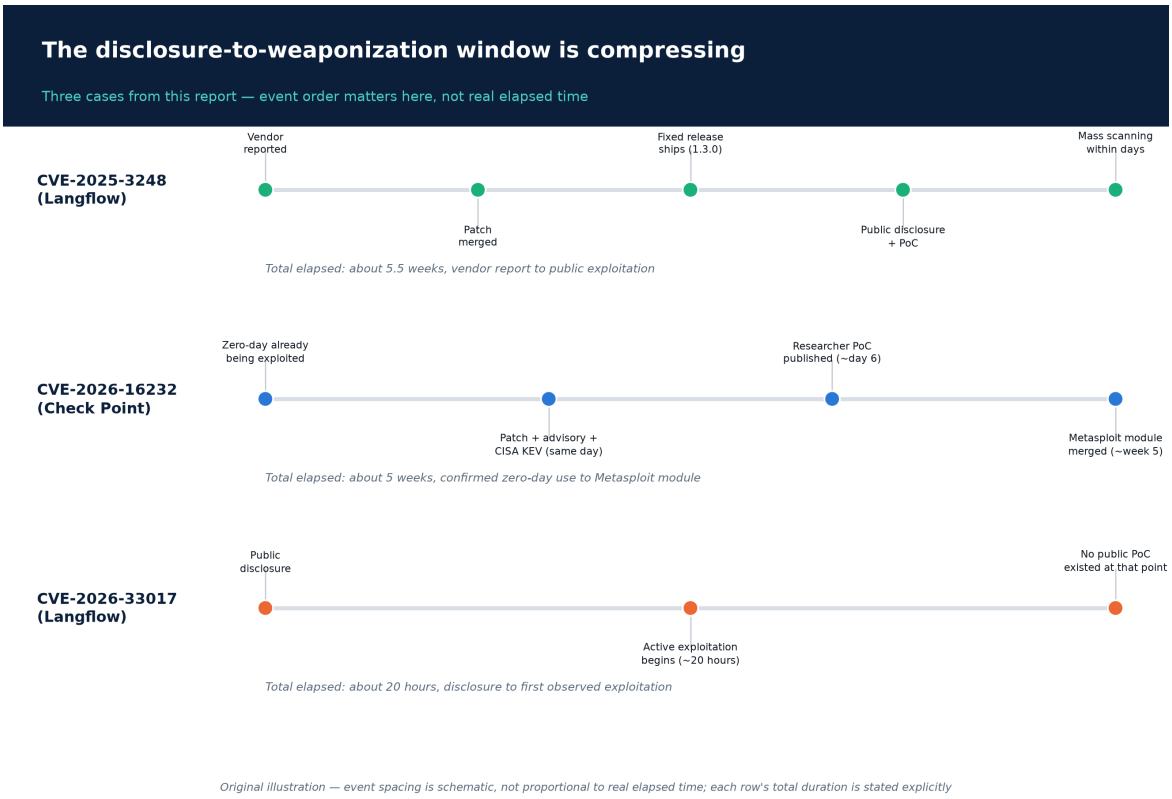

Figure D. Three cases from this report, each showing how little time now separates disclosure (or first exploitation) from a fully weaponized, public exploit path.

Original illustration, elapsed time is schematic to fit widely differing durations on one scale; each row's real total duration is stated explicitly.

7. Consolidated interim mitigations and hardening checklist

For teams that cannot patch one or more of these products immediately, in priority order:

Network exposure (applies to all four products)

- None of these products should have their administrative/API interface directly reachable from the public internet. Check Point's own root-cause statement for CVE-2026-16232 is that exploitation required direct internet exposure; the same is true of the Langflow/Flowise mass-scanning campaigns described above. Put every one of these consoles behind a VPN, bastion host, or private network segment with allow-listed administrator source IPs.
- Where a product-specific access-control list exists (Check Point's “Trusted Clients”), configure it to specific IPs/subnets rather than “Any.”

Check Point SmartConsole

Restrict Trusted Clients; firewall the management interface; verify implied control-connection rules; monitor audit logs for `Authentication method: application token`; patch to the JHA Take listed in §2.2 (or migrate off unsupported tracks).

Tenable Security Center

Restrict report-authoring and SCAP-audit-file-upload rights to trusted administrators; monitor the service account's process tree for anomalous spawning; patch to 6.9.0 (no interim hotfix exists).

Langflow

Disable `AUTO_LOGIN` unless required; monitor/alert on unauthenticated calls to `/api/v1/validate/code`, `/api/v1/build_public_tmp/*`, and `/api/v1/auto_login`; if the instance has ever been internet-exposed on a pre-1.3.0 build, treat it as a compromise-assessment candidate (Flodrix/Gafgyt/JADEPUFFER precedent), not just a patch target; patch to the latest release (1.10.1 at the time of writing) and independently re-verify the fix given JFrog's report of an incomplete patch for CVE-2026-33017.

Flowise

If the instance predates 3.0.6/3.1.2 and was ever internet-exposed, treat it as **presumptively compromised** per Cloud Security Alliance guidance and rotate every credential it held (LLM API keys, database credentials, cloud secrets) in addition to patching; enable built-in API authentication (disabled by default in older releases) and require it behind a reverse proxy; given the project's end-of-life status, begin migration planning now rather than treating the current release as a stable long-term target.

General AI-agent-framework hardening (Langflow, Flowise, and comparable platforms)

- Treat any “execute code” / “custom tool” / “CSV agent” / “code interpreter” feature as requiring the same network and privilege isolation as a remote shell, because that is functionally what it is.

- Run these platforms' code-execution components in a sandboxed, network-egress-restricted environment (container with no outbound internet access beyond explicitly allow-listed LLM API endpoints) so that even a successful code-injection exploit cannot reach internal systems or exfiltrate data trivially.
- Inventory and rotate any credentials (cloud, database, LLM API keys) stored in or accessible to these platforms on a schedule independent of known-incident status, given how frequently new CVEs have appeared in this product category.
- Subscribe to each vendor's/project's security-advisory feed directly (GitHub Security Advisories for Langflow/Flowise; Check Point's sk-article notifications; Tenable's TNS advisory mailing list) rather than relying solely on CISA KEV, since critical CVEs with public exploit code do not always appear in KEV.

8. Open questions

- **Exact CVSS/version data gaps.** For CVE-2026-55255 and CVE-2026-0770 (both confirmed on CISA KEV), this research could not independently confirm a published NVD CVSS score or a precise affected-version lower bound beyond the fix landing at Langflow 1.9.0. Defenders relying on risk-scoring these two should consult IBM's/Langflow's own advisory pages directly rather than the figures found via secondary aggregators.
- **Whether CVE-2025-34291's exploitation predates its public disclosure.** CrowdSec-reported exploitation beginning January 23, 2026 significantly predates the May 21, 2026 KEV addition; it is not established from public sources whether this reflects delayed public disclosure, delayed detection, or delayed KEV processing — worth monitoring for a fuller vendor timeline.
- **Whether Tenable Security Center's CVE-2026-19626/19681 will see confirmed exploitation.** As of this report, no ITW evidence exists, but the authenticated-access prerequisite, the existence of a public technical write-up, and the now-public Metasploit modules together create conditions historically associated with rapid follow-on exploitation in this report's other cases. This is a forward-looking risk to track via GreyNoise/Shadowserver reporting and Tenable's own advisory updates, not a settled fact.
- **The completeness of Langflow's 1.9.0/1.10.1 patches.** Given JFrog's documented finding that an earlier “fixed” Langflow release remained exploitable, defenders should treat each new Langflow patch claim with independent verification rather than assuming full remediation from the version number alone.
- **Post-archival security posture for Flowise.** It is not yet established whether any community fork will emerge as a de facto maintained successor, nor whether FlowiseAI will issue any further ad hoc security fixes despite the stated end-of-life. Organizations should treat this as an open, monitored question rather than assume either outcome.
- **Scale of Tenable Security Center's authenticated attack surface.** Because both weaponized Tenable CVEs require only a low-privileged, non-administrative account, the practical risk depends heavily on how broadly individual organizations have delegated report-creation and audit-file-upload rights — a fact this report cannot assess in aggregate and that each organization must audit locally.

Note on currency of this report: All Metasploit module data, CVE identifiers, and KEV entries reflect the state of public information as of August 30, 2026. Several vulnerabilities discussed here (particularly the Tenable Security Center pair and CVE-2026-56274 for Flowise) were disclosed and weaponized within the preceding two weeks; defenders should re-check CISA's KEV feed, NVD, and each vendor's advisory page directly before finalizing remediation priorities, since exploitation status can change rapidly for freshly weaponized CVEs.

Sources

Full source ledger for every claim in this report, in citation order. Each source name links directly to the primary document.

#	Claim	Evidence type	Source
1	Check Point SmartConsole module, authorship, CVE mapping	Primary — vendor blog	Rapid7, Metasploit Wrap-Up, Aug 28 2026
2	CVE-2026-16232 root cause, exploitation chain, figures	Primary — vendor technical research	Rapid7 Technical Analysis
3	CVE-2026-16232 exploited-in-the-wild confirmation, IOCs	Primary — vendor threat report	Rapid7 ETR Report
4	CVE-2026-16232 vendor advisory, affected/fixed versions, mitigations	Primary — vendor advisory	Check Point sk185169
5	CVE-2026-16232 active-exploitation statement, timeline, IOCs, mitigations	Primary — vendor blog	Check Point Security Advisory Blog
6	CVE-2026-16232 CVSS scores, CWE, dates	Primary — government database	NVD
7	CVE-2026-16232 KEV addition, due date	Primary — government catalog	CISA KEV JSON feed
8	CVE-2026-16232 Rapid7 PoC release coverage	Secondary — news	The Hacker News
9	CVE-2026-16232 exploited zero-day coverage	Secondary — news	BleepingComputer
10	CVE-2026-16232 exploited zero-day coverage	Secondary — news	Cybersecurity Dive
11	CVE-2026-16232 vulnerability scanner coverage	Secondary — vendor blog	Greenbone
12	Check Point Metasploit module source, PR	Primary — code repository	rapid7/metasploit-framework PR #21731
13	Check Point standalone PoC	Primary — researcher code repository	sfewer-r7/CVE-2026-16232
14	CISA BOD 26-04 details, risk-tiered remediation deadlines	Primary — government directive	CISA BOD 26-04
15	BOD 26-04 explainer	Secondary — vendor blog	Tenable
16	CVE-2026-19626, CVE-2026-19681 CVSS, CWE, description	Primary — government database	NVD (CVE-2026-19626)

#	Claim	Evidence type	Source
17	CVE-2026-19681 CVSS, CWE, description	Primary — government database	NVD (CVE-2026-19681)
18	Tenable Security Center affected/fixed versions, full CVE batch	Primary — vendor advisory	Tenable TNS-2026-22
19	CVE-2026-19626 vendor CVE page	Primary — vendor database	Tenable CVE page
20	Earlier, distinct Tenable advisory (context, not conflated)	Primary — vendor advisory	Tenable TNS-2026-19
21	TNS-2026-19 national CERT advisory	Primary — government advisory	Canadian Centre for Cyber Security AV26-724
22	Tenable Security Center Metasploit modules, PRs	Primary — code repository	rapid7/metasploit-framework PR #21820, #21802
23	Public discussion of eval-injection exploit for Tenable Security Center	Secondary — community forum	KSEC Forum
24	CVE-2025-3248 root cause (decorator/default-argument bypass)	Primary — researcher disclosure	Horizon3.ai
25	CVE-2025-3248 CVSS, CWE, affected versions	Primary — government database	NVD
26	CVE-2025-3248 GitHub Security Advisory	Primary — vendor/repo advisory	GHSA-rvqx-wpfh-mfx7
27	CVE-2025-3248 fix release	Primary — code repository	Langflow 1.3.0 release
28	CVE-2025-3248 Metasploit module	Primary — vendor module database	Rapid7 DB
29	CVE-2025-3248 Metasploit module PR	Primary — code repository	rapid7/metasploit-framework PR #20022
30	CVE-2025-3248 CISA KEV addition	Primary — government alert	CISA
31	CVE-2025-3248 KEV coverage	Secondary — news	The Hacker News
32	Flodrix botnet campaign exploiting CVE-2025-3248	Primary — vendor threat research	Trend Micro Research
33	Early exploitation-attempt telemetry for CVE-2025-3248	Primary — researcher diary (mirrored)	SANS ISC (via Harvard Tagteam mirror)
34	JADEPUFFER agentic ransomware using CVE-2025-3248	Primary — vendor threat research	Sysdig
35	JADEPUFFER coverage	Secondary — news	The Hacker News
36	JADEPUFFER coverage	Secondary — news	BleepingComputer
37	JADEPUFFER coverage	Secondary — news	SecurityWeek

#	Claim	Evidence type	Source
38	JADEPUFFER coverage	Secondary — news	CyberScoop
39	Langflow exposed-instance count estimate	Secondary — vendor/researcher blog	SaaS Sentinel
40	CVE-2026-27966 root cause, version, CVSS	Primary — government database	NVD
41	CVE-2026-27966 Metasploit module	Primary — code repository	rapid7/metasploit-framework PR #21260
42	CVE-2026-33017 root cause, exploitation timeline	Primary — vendor threat research	Sysdig
43	CVE-2026-33017 CISA KEV addition	Primary — government alert	CISA
44	CVE-2026-33017 Metasploit module	Primary — code repository	rapid7/metasploit-framework PR #21700
45	CVE-2026-33017 incomplete-patch finding	Primary — vendor research	JFrog Security Research
46	CVE-2025-34291 root cause, affected version	Primary — vendor/researcher blog	Obsidian Security
47	CVE-2025-34291 CISA KEV addition, exploitation timeline	Secondary — news	The Hacker News
48	CVE-2026-55255, CVE-2026-0770 KEV details	Primary — government catalog	CISA KEV JSON feed
49	CVE-2026-0770 fix reference	Primary — code repository	Langflow v1.9.0 release
50	CVE-2026-9198 root cause, module, disclosure date	Primary — code repository	rapid7/metasploit-framework module source
51	CVE-2026-9198 KEV addition	Primary — government catalog	CISA KEV JSON feed
52	CVE-2025-26319 root cause, CVSS	Primary — GitHub Security Advisory	GHSA-69jq-qr7w-j7qh
53	CVE-2025-8943 root cause, CVSS	Primary — government database / vendor research	NVD; JFrog JFSA-2025-001380578
54	CVE-2025-59528 root cause, CVSS 10.0	Primary — government database / vendor research	NVD; JFrog JFSA-2025-001379925
55	CVE-2026-41264 root cause, CVSS	Primary — researcher advisory	ZDI-26-364
56	CVE-2026-41264 GitHub Security Advisory	Primary — repo advisory	GHSA-3hjv-c53m-58jj

#	Claim	Evidence type	Source
57	CVE-2026-56274 root cause, CVSS (both versions)	Primary — government database	NVD
58	CVE-2026-56274 GitHub Security Advisory	Primary — repo advisory	GHSA-m99r-2hxc-cp3q
59	Flowise Metasploit modules, PRs	Primary — code repository	rapid7/metasploit-framework PR #20705, PR #21616
60	Flowise active exploitation (CVE-2025-59528, -8943, -26319), exposed-instance count	Secondary — news	The Hacker News
61	Flowise active exploitation coverage	Secondary — news	BleepingComputer
62	Flowise active exploitation coverage	Secondary — news	SecurityAffairs
63	Flowise “treat as compromised” guidance	Primary — research organization note	Cloud Security Alliance
64	Flowise IPS signatures for CVE-2025-59528	Primary — vendor blog	SonicWall
65	Flowise CVE-2024-31621 detection template	Primary — community detection ruleset	ProjectDiscovery nuclei-templates
66	Flowise project archival/sunset timeline	Primary — vendor announcement	FlowiseAI
67	Flowise sunset community discussion	Secondary — community forum	Hacker News
68	Rapid7 exploitation-flow and audit-log figures (Visual assets)	Primary — vendor technical research	Rapid7 Technical Analysis