CSAI Foundation | Cloud Security Alliance

Langflow CVE-2025-34291: RCE in AI Workflow Platforms

2026-05-22

AI-assisted Rapid Research

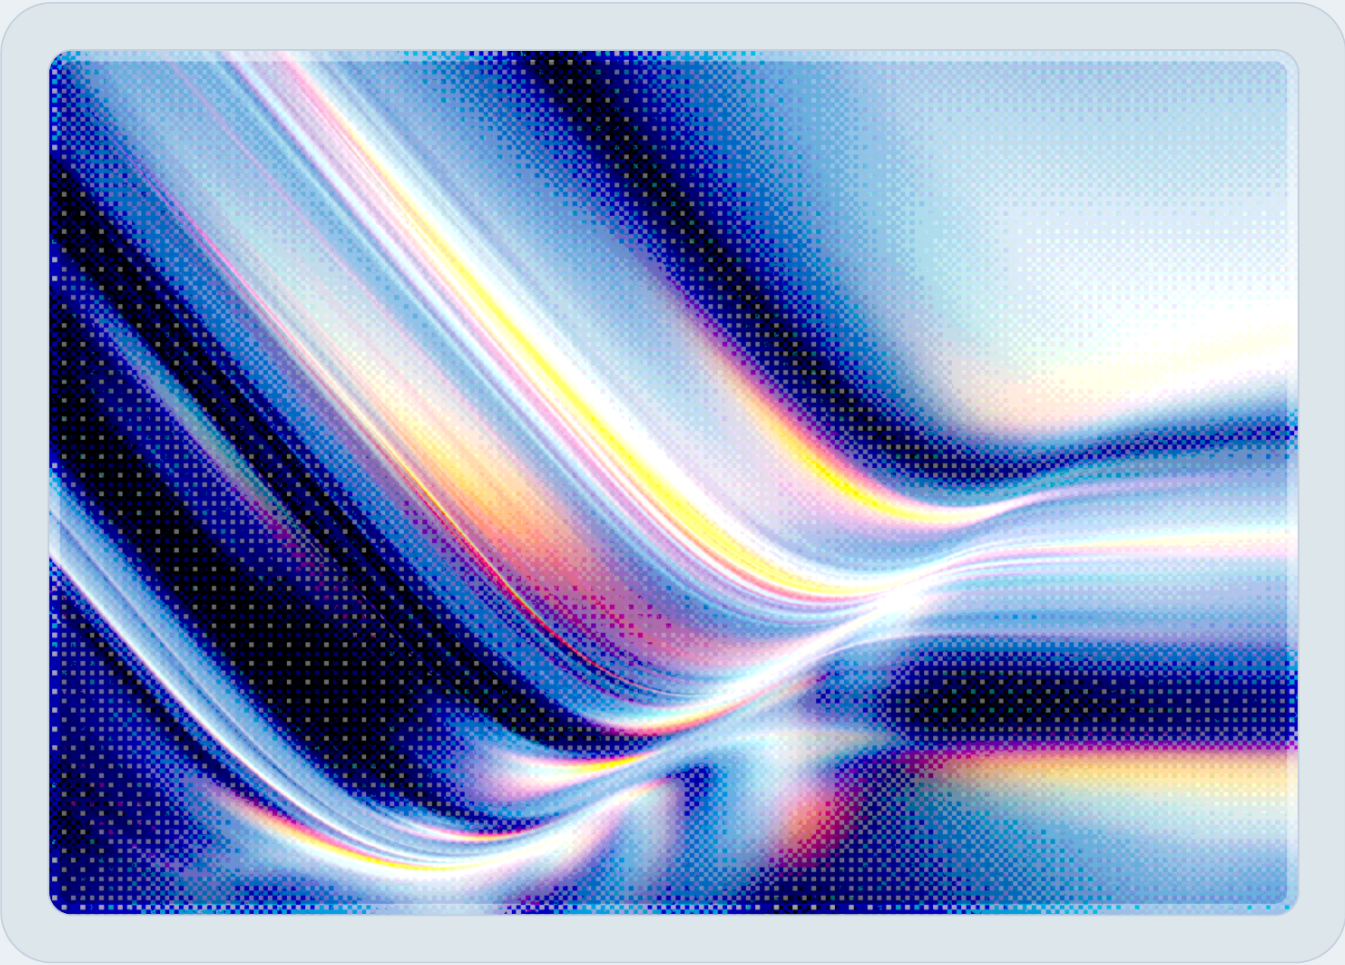

CSAI

CSA cloud security alliance®

Key Takeaways

- CVE-2025-34291 is a critical vulnerability chain (CVSS v4.0: 9.4) in Langflow versions up to and including 1.6.9, enabling full account takeover and remote code execution through a single victim browser visit to an attacker-controlled page [1].
- The exploit requires no stolen credentials: it chains a misconfigured CORS policy, a refresh token cookie scoped for cross-site delivery, and the absence of CSRF protection on the token refresh endpoint — culminating in access to authenticated code execution endpoints [2].
- CISA added CVE-2025-34291 to its Known Exploited Vulnerabilities (KEV) catalog on May 21, 2026, establishing a June 4, 2026 federal remediation deadline under Binding Operational Directive 22-01 [3].
- Active exploitation was first observed on January 23, 2026; an earlier vulnerability, CVE-2025-3248, enabled the Flodrix botnet to compromise publicly exposed Langflow instances beginning in mid-2025 [4, 5].
- Successful compromise exposes stored credentials within the Langflow workspace — including LLM provider API keys, third-party service tokens, and database credentials configured as Langflow integrations — extending the scope of potential compromise well beyond the orchestration platform itself [2].
- Organizations should immediately upgrade to Langflow 1.9.3 or later and restrict cross-origin credential delivery pending a full audit of any AI workflow tooling exposed to user-facing networks [3, 6].

Background

Langflow is an open-source visual platform for building multi-step AI agent pipelines. Developed and maintained with backing from DataStax, it allows engineers and data scientists to chain together large language model calls, retrieval-augmented generation components, tool integrations, and custom Python logic through a drag-and-drop interface [7]. Langflow appears to be among the more widely adopted tools in this category, reflecting a broader industry movement toward agentic AI architectures, where complex tasks are decomposed across networks of coordinated AI components that invoke external services, write to databases, and execute code in response to model outputs.

This last capability — executing arbitrary Python code as part of pipeline construction — is not an unintended side effect of Langflow's design but an intentional product feature. Users supply custom logic through a code validation endpoint that evaluates Python expressions to enable dynamic workflow construction. The power of this feature is precisely what makes it dangerous: any path to authenticated access on a Langflow instance effectively becomes a path to arbitrary code execution on the underlying host. The question, then, is not whether an authenticated code execution endpoint is dangerous in principle, but whether the authentication boundary protecting it is sound.

Langflow has accumulated a pattern of vulnerabilities centered on this boundary. CVE-2025-3248, disclosed in mid-2025, exposed the same code validation endpoint to entirely unauthenticated requests in versions prior to 1.3.0, carrying a CVSS score of 9.8 [8]. Threat actors rapidly weaponized public proof-of-concept exploits, deploying the Flodrix botnet — a DDoS capability descended from LeetHozer malware — through compromised Langflow instances exposed on public-facing infrastructure [5]. CISA subsequently added CVE-2025-3248 to its KEV catalog [15]. A follow-on vulnerability, CVE-2026-33017, disclosed in March 2026, returned to the same `exec()`-based validation function and was, according to Barrack AI, weaponized within twenty hours of the advisory's publication, with attackers harvesting API keys for OpenAI, Anthropic, and AWS from compromised instances [9]. Against this backdrop, CVE-2025-34291 represents a third major exploitation vector against Langflow's code execution surface — this time achieved through web application security flaws rather than missing authentication checks.

Security Analysis

The Exploit Chain

CVE-2025-34291 was identified and disclosed by researchers at Obsidian Security. Its severity stems not from a single coding error but from the compounding effect of three separate misconfigurations that, in combination, allow a fully unauthenticated attacker to hijack an authenticated session and invoke code execution — all triggered by a victim visiting a malicious webpage [1, 2].

The first element is Langflow's CORS configuration. In affected versions, the platform sets `allow_origins='*'` while simultaneously enabling `allow_credentials=True`. This combination is explicitly disallowed by the CORS specification, which requires that credential-bearing cross-origin requests be explicitly scoped to trusted origins rather than wildcards. Langflow's CORS middleware, despite the `allow_origins='*'` configuration setting, processes cross-origin requests by echoing the requester's `Origin` header in its response — effectively granting any domain

the specific-origin credential access the CORS specification reserves for explicitly trusted origins. An attacker operating from any domain can therefore issue cross-origin requests to the Langflow server that carry the victim's browser cookies and receive the full response [2].

The second element is the configuration of the `refresh_token_lf` session cookie. This cookie is set with `SameSite=None; Secure`, which, by specification, allows it to accompany cross-site requests over HTTPS. It carries a validity window of up to one week, providing a substantial exploitation window once a victim has authenticated to a Langflow instance [2]. The combination of the permissive CORS setting and the cross-site cookie means that a script executing in the attacker's browser page context can reach the Langflow server with the victim's refresh token attached.

The third element is the absence of CSRF protection on the `/api/v1/refresh` endpoint. This endpoint issues a fresh pair of access and refresh tokens, but it relies solely on the presence of the cookie for authentication — implementing no additional verification that the request originated from a trusted context [2]. An attacker-controlled page can therefore call this endpoint cross-origin, supply the victim's cookie automatically, and receive in response a fresh access token that is valid for all authenticated Langflow API operations.

With a valid access token in hand, the attacker calls authenticated endpoints including the code validation API, invoking arbitrary Python on the underlying server. The attack requires no prior knowledge of credentials, no phishing of passwords, and no malware installation on the victim's machine. The victim's single act of visiting a malicious URL while holding an active Langflow session is sufficient for full compromise.

Scope of Compromise

The consequences of a successful CVE-2025-34291 exploit extend beyond the Langflow application layer. Langflow deployments typically store connection credentials for every external service integrated into user workflows: API keys for OpenAI, Anthropic, and other LLM providers; database connection strings; vector store credentials; and webhook secrets for downstream services [2, 9]. An attacker who achieves code execution on the Langflow host can enumerate and exfiltrate all of these credentials, converting a single browser-based session hijack into persistent access across the organization's broader AI infrastructure.

This attack surface geometry is characteristic of AI orchestration platforms more generally. Because their core function is to serve as trusted intermediaries between model APIs and enterprise data systems, they accumulate high-value credentials by design. Security teams that have not explicitly inventoried the credential stores embedded in AI workflow tooling may therefore find that a Langflow compromise yields lateral access to resources far removed from the workflow platform itself.

Active Exploitation and Threat Actor Context

CrowdSec researchers documented active exploitation of CVE-2025-34291 beginning January 23, 2026, several months before CISA's formal KEV designation [4]. The exploitation pattern shares tactical overlap with the CVE-2025-3248 campaigns: internet-exposed Langflow instances are identified through scanning services such as Shodan, and public or reconstructed proof-of-concept code is used to trigger compromise [5]. The Flodrix botnet — the most extensively documented threat actor cluster in reporting on Langflow attacks, as analyzed by Trend Micro during the CVE-2025-3248 campaign — represents a recurring presence in campaigns targeting internet-facing AI tooling, though the broader attacker population is not limited to a single campaign [5, 10]. Separately, reporting on the CISA KEV designation indicates that MuddyWater, an Iranian state-sponsored group, leveraged CVE-2025-34291 for initial network access, confirming that both criminal and nation-state actors have incorporated the vulnerability into active operations [10].

CISA's decision to add CVE-2025-34291 to the KEV catalog on May 21, 2026, reflects an assessment that exploitation is not speculative but ongoing [3]. A 14-day remediation deadline establishes the June 4, 2026 compliance horizon for federal civilian agencies, and CISA has urged private-sector organizations to treat KEV designations as a prioritization signal regardless of their regulatory status.

Architectural Root Cause

Viewed at the platform level, the recurring vulnerability pattern in Langflow reflects a design tension common across AI workflow orchestration tools: providing powerful built-in code execution for legitimate customization while maintaining an authentication boundary robust enough to prevent that same execution capability from being weaponized. Both CVE-2025-3248 and CVE-2025-34291 attacked different layers of this boundary — the former bypassing authentication entirely, the latter subverting it through web application security flaws. The recurrence suggests that the boundary has not been treated with the same rigor as the execution capability itself.

Langflow is not likely unique in this regard — the design tension between powerful code execution and robust authentication boundaries appears across comparable agentic AI platforms. Similar visual pipeline tools and AI agent frameworks that offer built-in Python execution, shell access, or dynamic tool invocation face structurally analogous risks whenever their authentication and request validation logic is imperfect. The same analysis applies to any agentic AI orchestration layer that stores credentials and executes code on behalf of its users.

Recommendations

Immediate Actions

Organizations running Langflow should treat this as an emergency patching event. The vendor has addressed CVE-2025-34291 in Langflow version 1.9.3, and all deployments running versions up to and including 1.6.9 should be upgraded without delay [6]. For deployments where immediate patching is not operationally feasible, two compensating controls are available at the configuration layer: setting the `LANGFLOW_CORS_ALLOW_CREDENTIALS` environment variable to `False` disables credential delivery in cross-origin contexts, directly severing the exploit chain; alternatively, setting `LANGFLOW_CORS_ORIGINS` to an explicit allowlist of trusted origins restricts which domains may issue credentialed requests [2]. Either control should be treated as temporary — complete remediation requires upgrading to a patched release.

All API keys, LLM provider tokens, and service credentials stored in any Langflow instance that may have been exposed to untrusted network traffic since January 2026 should be treated as potentially compromised and rotated. This includes credentials for OpenAI, Anthropic, AWS, vector databases, and any other services configured as Langflow integrations.

Short-Term Mitigations

Security teams should conduct an inventory of all AI workflow and orchestration tooling deployed within their environments, identifying any instances accessible from user-facing networks, public cloud endpoints, or insufficiently segmented development environments. AI pipeline tools should operate behind authentication perimeters, and direct exposure of orchestration APIs to the public internet should be treated as an elevated-risk configuration requiring explicit justification. Web application firewalls and reverse proxies placed in front of Langflow deployments can enforce additional origin validation and rate limiting as layered controls while patch deployment proceeds.

Log review should focus on anomalous calls to the `/api/v1/refresh` endpoint — particularly requests originating from unexpected origins or IP addresses — and any subsequent authenticated API calls that differ from baseline workflow automation patterns. Given that the Flodrix botnet has actively targeted Langflow, indicators associated with that campaign, including C2 communication patterns and lateral DDoS traffic, should be incorporated into detection queries.

Strategic Considerations

The repeated exploitation of Langflow's code execution surface points to a broader governance gap that organizations deploying AI orchestration tooling should address at the architectural level. Any platform that combines stored credentials with arbitrary code execution capabilities warrants the same threat-modeling discipline applied to privileged infrastructure: network isolation, continuous vulnerability monitoring, secret rotation policies, and code review requirements for custom component logic. These characteristics create favorable conditions for adversaries seeking to convert a single web application flaw into mass credential harvesting, and security programs that have not explicitly accounted for AI orchestration platforms within their risk models should do so.

Organizations procuring AI workflow platforms should evaluate vendors' security disclosure and patching track records as part of procurement criteria. Langflow's three-CVE sequence — spanning unauthenticated RCE, session hijack via web application misconfiguration, and a second unauthenticated RCE — represents a pattern that security architects should weigh when selecting foundational components for agentic AI infrastructure. Preference should be given to platforms that enforce a strict separation between workflow orchestration and code execution, minimize credential persistence, and provide audit logging granular enough to detect anomalous API usage. Security teams conducting red team exercises on AI infrastructure should explicitly include AI workflow orchestration platforms in scope, as these components may be underrepresented in adversarial testing given their relatively recent adoption and the tendency to treat AI infrastructure as application-layer rather than platform-layer risk.

CSA Resource Alignment

The vulnerabilities documented in CVE-2025-34291 align with threat surfaces analyzed in CSA's MAESTRO framework, which provides a seven-layer model for threat modeling agentic AI systems [11]. The orchestration layer — where Langflow operates — is identified in MAESTRO as a high-risk integration point because it aggregates credentials, issues commands to downstream services, and frequently executes code in response to model-generated instructions. The MAESTRO framework addresses this exposure through its defense-in-depth and secure inter-layer communication guidance, and organizations can use the framework to systematically assess how their Langflow deployment controls compare against the orchestration-layer security model it defines.

CSA's AI Controls Matrix (AICM) provides governance-level controls directly applicable to this class of vulnerability [12]. Relevant AICM domains include access management for AI service integrations, secrets management requirements for pipeline credentials, and vulnerability management expectations for AI

infrastructure components. The AICM extends the Cloud Controls Matrix (CCM) to address AI-specific risks, and its controls can serve as a structured baseline for evaluating Langflow deployment configurations against organizational security policy.

The CSA Agentic AI Red Teaming Guide, developed by the AI Organizational Responsibilities Working Group, addresses testing methodologies for agentic systems including orchestration layers, and should inform any penetration testing or security assessment of Langflow or similar platforms [13]. The guide's treatment of code injection against agent tool-call interfaces is particularly relevant to Langflow's execution surface. Security teams conducting red team exercises on AI infrastructure should explicitly include AI workflow orchestration platforms in scope, recognizing that these components have historically received less adversarial scrutiny than their actual attack surface warrants.

CSA's Zero Trust guidance reinforces the architectural principle that Langflow deployments should not be treated as implicitly trusted by adjacent services simply because they operate within a cloud environment or corporate network. Every service integration connected to a Langflow instance should require explicit authentication from that instance, and the credentials used should be scoped to the minimum permissions required for each workflow's operation [14].

References

[1] Obsidian Security. "CVE-2025-34291: Critical Account Takeover and RCE Vulnerability in the Langflow AI Agent & Workflow Platform." Obsidian Security, 2026.

[2] GitHub Security Advisories. "GHSA-577h-p2hh-v4mv: Langflow versions up to and including 1.6.9 origin validation error." GitHub, 2026.

[3] CISA. "CISA Adds Two Known Exploited Vulnerabilities to Catalog." CISA, May 21, 2026.

[4] CrowdSec. "CVE-2025-34291 Exploited in the Wild: LangFlow AI Framework Under Fire." CrowdSec VulnTracking, 2026.

[5] Trend Micro. "Critical Langflow Vulnerability (CVE-2025-3248) Actively Exploited to Deliver Flodrix Botnet." Trend Micro Research, 2025.

[6] NIST National Vulnerability Database. "CVE-2025-34291 Detail." NVD, 2026.

[7] Langflow. "Langflow: The Open-Source AI Application Platform." DataStax, 2026.

[8] Keysight Technologies. "CVE-2025-3248: Unauthenticated Remote Code Execution in Langflow." Keysight Blog, June 2025.

[9] Barrack AI. "Langflow Got Hacked Twice Through the Same exec() Call." Barrack AI Blog, 2026.

[10] The Hacker News. "CISA Adds Exploited Langflow and Trend Micro Apex One Vulnerabilities to KEV." The Hacker News, May 2026.

[11] Cloud Security Alliance. "Agentic AI Threat Modeling Framework: MAESTRO." CSA Blog, February 2025.

[12] Cloud Security Alliance. "AI Controls Matrix." CSA AI Controls Matrix Working Group, 2025.

[13] Cloud Security Alliance. "Agentic AI Red Teaming Guide." CSA AI Organizational Responsibilities Working Group, 2025.

[14] Cloud Security Alliance. "Zero Trust Working Group." CSA Zero Trust Working Group.

[15] CISA. "Known Exploited Vulnerabilities Catalog." CISA, 2025.

