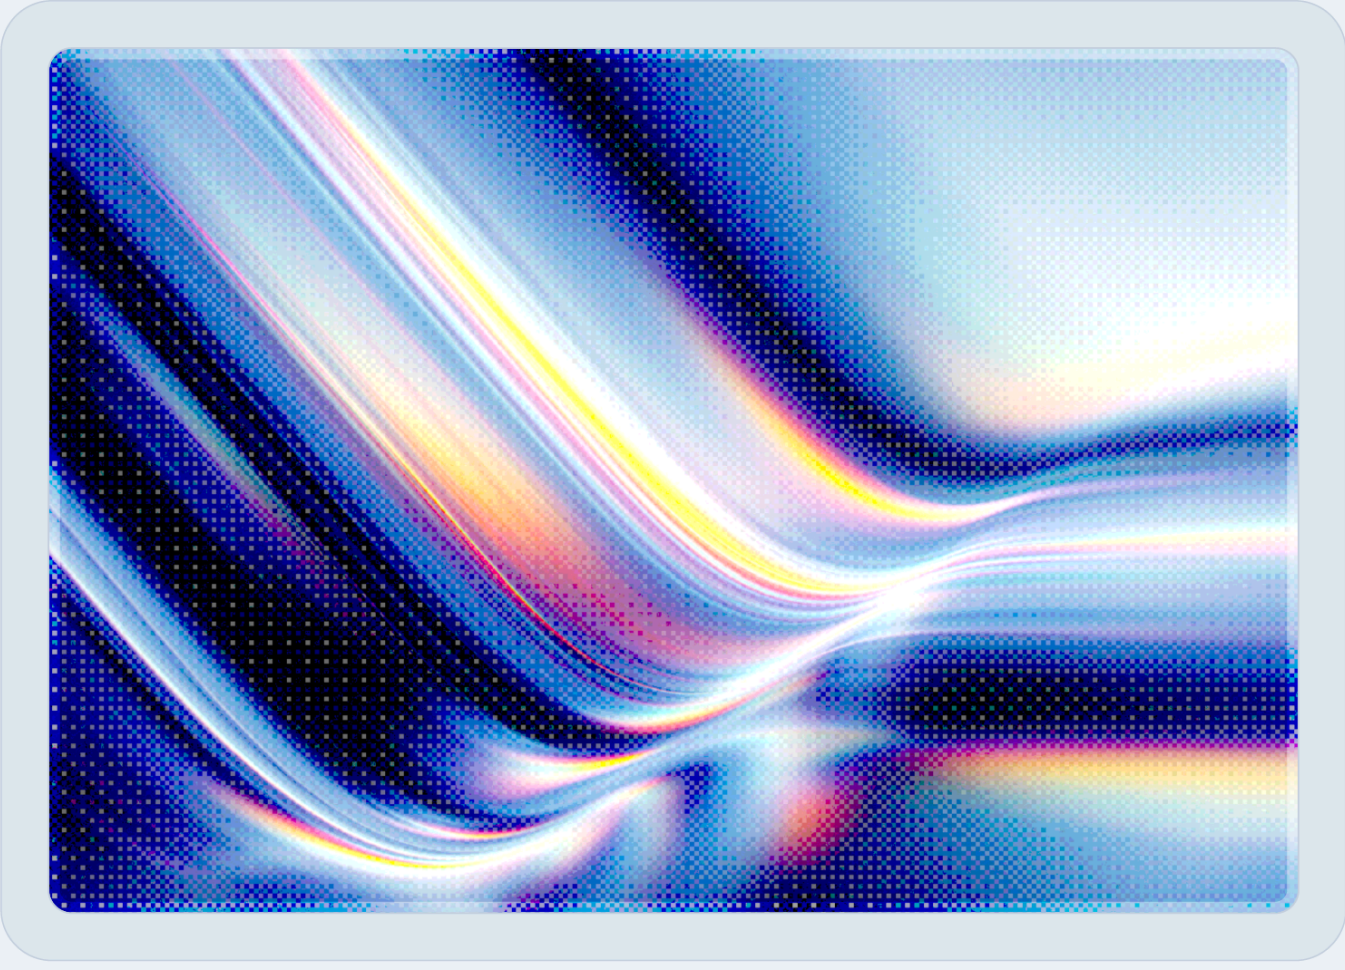
CSAI Foundation | Cloud Security Alliance

Langflow 1.13.0: Two Unauthenticated Critical Flaws and an Unpatched 1.12 Branch

CVE-2026-93680 and CVE-2026-110186: Security Implications and Guidance

2026-11-03

AI-assisted Rapid Research



CSAI

CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

IBM, which now acts as the CVE numbering authority for Langflow, published a security bulletin on October 28, 2026 disclosing two critical vulnerabilities in Langflow OSS that require no account to exploit [1][8]. CVE-2026-93680 is an authentication bypass, and CVE-2026-110186 is an arbitrary code execution flaw arising from improper control of code generation. Both carry an IBM-assigned CVSS 3.1 base score of 9.8, and both affect every release from 1.0.0 through 1.12.8 [1]. The only remediation IBM names is Langflow 1.13.0, and IBM lists no workarounds or mitigations for organizations that cannot upgrade [1][3].

The disclosure has an unusual shape that defenders should understand. Langflow 1.12.8 and 1.13.0 were published to PyPI on the same day, October 22, 2026, and 1.12.8 remains within the affected range [1][9]. As of November 3, 2026, no later 1.12.x release exists on PyPI, and Langflow has not stated whether one is planned, so organizations that stayed on the 1.12 line currently have no patched release in that branch and must move to 1.13.0 [4][9]. That upgrade turns automatic login off by default, and in-place upgrades from deployments that relied on auto-login need deliberate preparation before the server will start or accept sign-ins [4][8].

Public tooling already exists. Rapid7's October 23 Metasploit release, which named CVE-2026-93680 five days before IBM's bulletin, added an auxiliary module that opens a Langflow MCP session without credentials and enumerates the server's tools; the module does not itself execute code [5][6]. Neither critical CVE was in CISA's Known Exploited Vulnerabilities (KEV) catalog as of its October 30 update, but six earlier Langflow CVEs are, at least three of which KEV describes as unauthenticated code execution flaws [10]. Given that history, CSA recommends treating any internet-reachable Langflow instance at or below 1.12.8 as an emergency patching priority.

Background

Langflow is an open-source, low-code platform for visually composing LLM applications, retrieval pipelines, and AI agents. Because Langflow's core function is to assemble and run Python components on behalf of users, a server that hosts it typically holds credentials for model providers, vector databases, and downstream business systems, and it can expose its flows as tools to Model Context Protocol (MCP) clients [6]. Those properties, together with the platform's KEV history described below, help explain its

appeal to attackers: a compromise yields both code execution on the host and a ready-made store of API keys and data connections. Langflow vulnerability reports are now handled by IBM following its acquisition of DataStax, Langflow's prior steward [6][8].

Langflow has an extensive record of exploited vulnerabilities, and CISA's KEV catalog currently lists six Langflow CVEs [10]. Four of them are especially relevant to this note: CVE-2025-3248, the unauthenticated code-validation endpoint flaw added in May 2025; CVE-2026-33017, a code injection flaw in the public flow build endpoint reachable without authentication; CVE-2026-55255, an authorization bypass; and CVE-2026-9198, a chain in which the `/api/v1/auto_login` endpoint issued superuser tokens to unauthenticated callers, who could then reach the code-validation endpoint [8][10][11][12][17]. The remaining two KEV entries, CVE-2025-34291 and CVE-2026-0770, also enable code execution [10]. Outside the KEV list, CSA has also documented active exploitation of CVE-2026-5027, a path traversal flaw leading to unauthenticated remote code execution [18]. CSA has published research notes on several of these issues, and a recurring theme across them is that development-convenience features, auto-login in particular, become network attack surface when Langflow is deployed beyond a single developer's workstation [11][12].

The current disclosure arrives within a much larger batch. IBM's October 19 bulletins addressed vulnerabilities fixed in 1.12.4 and 1.12.5, and on October 29 IBM published CVE records for 29 additional Langflow flaws, naming 1.12.4 and 1.12.5 as fixes [2][6][16]. One of those earlier-fixed flaws, CVE-2026-104339, is also an unauthenticated code execution issue rated 9.8 [6]. One independent tally counts 33 Langflow CVEs across the October bulletins, with every release from 1.0.0 through 1.12.8 affected by at least one of them and only 1.13.0 addressing all 33 [8]. The two critical flaws this note focuses on are distinguished by combining unauthenticated, network-reachable exploitation with the absence of any 1.12.x fix.

The following timeline summarizes the sequence of releases, bulletins, and public tooling relevant to this note.

Date (2026)	Event	Source
September 29	Langflow 1.12.4 released	[9]
October 6	Langflow 1.12.5 released	[9]
October 19	IBM bulletins: 1.0.0–1.12.3 affected, fixed in 1.12.4; 1.0.0–1.12.4 affected, fixed in 1.12.5	[2][16]

Date (2026)	Event	Source
October 22	Langflow 1.12.8 and 1.13.0 both released; 1.13.0 disables auto-login by default	[4][9]
October 23	Rapid7 Metasploit wrap-up adds auxiliary module naming CVE-2026-93680	[5]
October 28	IBM bulletin: CVE-2026-93680, CVE-2026-110186, and others; 1.0.0–1.12.8 affected, fixed in 1.13.0	[1]
October 29	IBM bulletin updated; CVE records for 29 further Langflow flaws published	[1][6]
October 30	CISA KEV update; neither critical CVE listed	[6][10]
November 3	Uganda CERT issues advisory on both critical CVEs	[7]

Security Analysis

The Two Unauthenticated Critical Flaws

IBM's bulletin is terse, and defenders should be careful to distinguish what it states from what commentators have inferred. IBM describes CVE-2026-93680 as permitting a remote attacker to "bypass security restrictions" due to improper authentication (CWE-287), and CVE-2026-110186 as allowing arbitrary code execution through improper control of code generation (CWE-94) [1]. Both carry the vector AV:N/AC:L/PR:N/UI:N, meaning network-reachable, low complexity, and requiring neither privileges nor user interaction [1]. The bulletin's summary, also reported by The Hacker News, characterizes one flaw as "an incomplete remediation of an auto-login security control" and another as "unauthenticated remote code execution exploitable through model context protocol tool enumeration" [1][6]. IBM does not publish endpoints, proofs of concept, or an explicit mapping of those summary descriptions to CVE identifiers [6][8].

The same bulletin lists two lower-severity Langflow issues that share the 1.12.8 ceiling. CVE-2026-110193 (CVSS 8.8) allows a low-privileged authenticated user to induce the Agentic Assistant to generate or submit component code using unblocked patterns, resulting in operating system command execution; CVE-2026-110192 (CVSS 6.5) involves insufficiently protected credentials that an authenticated attacker

can obtain [1]. The bulletin also lists three vulnerabilities in the Hono web framework (CVE-2026-84363, CVE-2026-84364, CVE-2026-84365), fixed upstream in Hono 4.13.5, although it does not state how those fixes map to Langflow's release [1].

The table below summarizes the Langflow-specific entries in IBM's October 28 bulletin, with the privileges column reflecting each entry's CVSS PR metric.

CVE	CVSS 3.1	Privileges Required	IBM Description (paraphrased)	Affected	Fixed
CVE-2026-93680	9.8	None	Authentication bypass via improper authentication	1.0.0–1.12.8	1.13.0
CVE-2026-110186	9.8	None	Arbitrary code execution via improper control of code generation	1.0.0–1.12.8	1.13.0
CVE-2026-110193	8.8	Low	Agentic Assistant generates component code that runs OS commands	1.0.0–1.12.8	1.13.0
CVE-2026-110192	6.5	Low	Sensitive information disclosure via insufficiently protected credentials	1.0.0–1.12.8	1.13.0

The Auto-Login and MCP Connection

The public evidence available so far points toward Langflow's auto-login feature and its MCP server interface as the relevant attack surface, though IBM has not confirmed the exact mechanics. Rapid7's October 23 release identified its new module by CVE-2026-93680, which shows the identifier was public in Metasploit five days before IBM's bulletin [5]. The module connects to Langflow's MCP endpoint at `/api/v1/mcp/sse` without credentials, establishes a session, and lists the tools the server offers [6][8]. Its author, Rapid7 contributor Richard Howe, tested it against Langflow 1.8.4 with automatic login enabled and lists 1.5.0 through 1.11.4 as affected, a narrower range than IBM's 1.0.0 through 1.12.8 [5][6]. The same Rapid7 release added three unauthenticated Langflow code execution modules for other CVEs (CVE-2026-0768, CVE-2026-0769, and CVE-2026-10134), none of which Rapid7 lists as affecting 1.12.x [5].

When `LANGFLOW_AUTO_LOGIN` is enabled, Langflow signs callers in as superusers without prompting, and its own code description warns that the setting "bypasses authentication and should only be used in development environments" [6]. That setting was still on by default in 1.12.8 [6]. It is reasonable to infer, though not confirmed by IBM, that deployments running with auto-login enabled are the most exposed population, and that the "incomplete remediation" language refers to protections added after the earlier CVE-2026-9198 auto-login chain that did not cover every route, such as MCP [6][11]. IBM's bulletin does not state whether servers with auto-login disabled are vulnerable, so organizations should not assume that disabling it on 1.12.x is a sufficient control [6].

The pairing of an authentication bypass with a code-generation flaw is the central reason for concern. Enumerating MCP tools without credentials is, by itself, an information disclosure. If, as IBM's summary language suggests, that same unauthenticated MCP path can be driven to code execution, the combined effect resembles the two-request chains seen in CVE-2026-9198 and CVE-2026-33017, both of which reached KEV status after in-the-wild exploitation [10][11][17]. Langflow's history also shows that attackers have targeted its flaws quickly; in the case of CVE-2026-9198, CSA documented exploitation attempts beginning even before public disclosure [11].

Why the 1.12 Branch Matters

Organizations commonly pin dependencies to a minor line and apply patch releases within it, a sensible practice for a platform whose minor versions can change storage backends and defaults. That practice does not provide protection in this case. PyPI shows 1.12.8 published at 17:08 UTC and 1.13.0 at 21:20 UTC on October 22, 2026; as of November 3, no 1.12.9 has appeared, while release candidates for 1.13.1 have [9]. The Hacker News likewise reported that PyPI listed no later 1.12 release as of November 2 [6]. An operator who diligently upgraded to the newest 1.12.x patch on October 22 therefore installed a release that IBM's October 28 bulletin lists as vulnerable to both critical flaws [1][9]. The Langflow 1.13.0 release notes make no explicit statement about continued support or an end-of-life date for 1.12.x, so whether a 1.12.x fix will follow remains an open question [4].

The move to 1.13.0 is not a drop-in patch. The release disables automatic login by default across official packages and Docker images, replaces local Chroma knowledge bases with SQLite, makes usernames case-insensitively unique, and makes timestamps timezone-aware [4]. Independent upgrade guidance reports that an in-place upgrade from an auto-login deployment fails to start with "Username and password must be set" until a superuser password is configured, and that because such instances already contain a `langflow` account with a generated password, setting `LANGFLOW_SUPERUSER_PASSWORD` alone yields "Incorrect username or password" [8]. The Langflow project added recovery guidance for this scenario in the 1.13.0 release [4]. The Hacker News

also cautions against deleting or reinstalling Langflow to regain access, since doing so can permanently remove flows [6]. These frictions are a plausible reason some operators will defer the upgrade, which lengthens the exposure window.

Exposure and Exploitation Status

As of the October 30 KEV update, neither CVE-2026-93680 nor CVE-2026-110186 had been added to the catalog, and the sources reviewed for this note report no confirmed in-the-wild exploitation [6][8] [10]. Absence from KEV is not evidence of absence of exploitation, however. Uganda's national CERT rated the issue Red on November 3, noting that a public Metasploit module already targets the authentication bypass and that AI workflow infrastructure is frequently deployed without the review applied to traditional production systems [7]. The 1.13.0 release also includes changes that bear on post-compromise impact, including an update to the `migrate_secret_key` script that rotates API key and MCP server secrets during secret-key migration, and a credential egress policy across provider components [4]. Separately, HOL's upgrade guidance notes that CVE-2026-77817 concerned exports that did not strip secrets, and that upgrading fixes only new exports; older exported flows may still contain live keys [8].

Recommendations

Immediate Actions

Organizations should first inventory every Langflow deployment, including developer workstations, notebooks, and container images, and determine its version using `pip show langflow`, the `/api/v1/version` endpoint, or the image tag [8]. Any instance at 1.12.8 or earlier should be upgraded to Langflow 1.13.0 or a later stable 1.13.x release, pinned to an explicit version, for example `pip install --upgrade "langflow==1.13.0"` or the `langflowai/langflow:1.13.0` image rather than a floating `latest` tag [1][8]. Before upgrading, teams should back up the database, secret key, and uploaded files, and if the deployment used auto-login, prepare the existing account for password sign-in following Langflow's installation and Docker upgrade instructions [1][13].

For any instance that cannot be upgraded within hours, the priority is to remove it from direct network reachability. Binding Langflow to `127.0.0.1` or placing it behind an authenticating reverse proxy that covers every path, including `/api/` and `/api/v1/mcp/`, are the most dependable interim controls

[8]. A simple exposure check is whether `/api/v1/mcp/sse` returns an event stream to an unauthenticated request rather than a 401 or 403 response [8]. Because IBM lists no official workaround, these steps reduce exposure but should not be treated as remediation [1].

Short-Term Mitigations

Over the following days, security teams should assume that any Langflow instance which was internet-reachable at an affected version may have been accessed, and should rotate every credential stored in it, including model provider keys, database connection strings, and MCP server secrets. Teams should also review web and application logs for unauthenticated requests to MCP endpoints, auto-login endpoints, and code-validation or flow-build routes, and for unexpected child processes spawned by the Langflow service. Where 1.12.x must run temporarily, operators should set `LANGFLOW_AUTO_LOGIN=false` with a superuser password and set `LANGFLOW_ENABLE_SIGNUP=false` [8]. For MCP stdio deployments, IBM's October 19 bulletin for the 1.12.5 fixes recommends `LANGFLOW_MCP_SERVER_INTERPRETER_HARDENING=true` and `LANGFLOW_CUSTOM_COMPONENT_ADMIN_ONLY=true`, while stating that these controls do not replace the upgrade [16]. None of these settings is confirmed to block CVE-2026-93680 or CVE-2026-110186, so network isolation remains necessary. Finally, previously exported flow files should be audited for embedded secrets, and any keys found in them should be rotated [8].

Strategic Considerations

This disclosure reinforces a lesson from Langflow's prior KEV entries: platforms whose purpose is to execute user-composed code should not depend on a single authentication layer, and development-oriented defaults like auto-login should never reach shared or production environments. Organizations should run Langflow in an isolated network segment, without root privileges, with outbound egress restricted to the model and data endpoints its flows actually need, so that a code execution flaw does not automatically become access to the broader environment. Version-pinning policies should also account for the possibility that a vendor ships critical fixes only in a newer minor line. In this case, no 1.12.x fix had been published as of November 3, and Langflow had not stated whether one is planned; subscribing to IBM's Langflow bulletins and monitoring PyPI for branch activity provides earlier warning of such situations. Finally, MCP servers exposed by agent platforms deserve the same authentication, inventory, and monitoring treatment as any other API, because tool enumeration endpoints provide attackers with reconnaissance and, as this case suggests, a potential path to execution.

CSA Resource Alignment

CSA's research note *Langflow Auto-Login Bypass Chains to Unauthenticated RCE* is the most directly relevant prior work [11]. It analyzed CVE-2026-9198, in which the auto-login endpoint issued superuser tokens to unauthenticated callers who then reached an `exec()`-based code-validation endpoint, and it recommended disabling auto-login outside development. IBM's bulletin summary describes one of the current flaws as an incomplete remediation of an auto-login security control, but IBM has not stated which CVE that description refers to [1][8]. If it is CVE-2026-93680, the attack surface CSA analyzed in August may not have been fully closed in the 1.12 line. Langflow 1.13.0's decision to disable auto-login by default is consistent with the recommendation made in that note, though IBM has not stated whether it is the operative fix.

CSA's note *Langflow Authorization Bypass Added to CISA's KEV Catalog* examined CVE-2026-55255 and its chaining with CVE-2026-33017, and it flagged running Langflow as root, internet exposure, default auto-login, and incomplete initial patches as recurring pitfalls [12]. Several of those pitfalls, namely internet exposure, default auto-login, and possibly incomplete initial patches, appear relevant to the present disclosure. The deployment hardening guidance in that note, including running Langflow without root privileges, applies directly to organizations that must keep 1.12.x running briefly. CSA's dedicated analysis of *CVE-2026-33017* and its note on the actively exploited *CVE-2026-5027 path traversal* provide further context on how unauthenticated Langflow flaws have been used against AI pipelines [17] [18].

At the control level, the AI Controls Matrix (AICM) v1.1 provides the framework for the recommendations above [14]. Its Threat and Vulnerability Management domain covers timely remediation and asset inventory for AI tooling, its Application and Interface Security domain addresses the authentication of API and MCP endpoints, and its Identity and Access Management domain addresses the elimination of default superuser access. For organizations deploying Langflow as agent infrastructure, MAESTRO, CSA's agentic AI threat modeling framework, offers a structured way to model how an unauthenticated tool-enumeration interface on an agent framework layer propagates risk to the data and deployment layers beneath it [15].

References

[1] IBM. "Security Bulletin: Langflow OSS is affected by multiple vulnerabilities." IBM Support, October 28, 2026 (updated October 29, 2026).

[2] IBM. "Security Bulletin: Langflow OSS is affected by multiple vulnerabilities (1.0.0–1.12.3, fixed in 1.12.4)." IBM Support, October 19, 2026.

[3] AiCybr. "Langflow OSS: Two Critical 9.8 Flaws Fixed in 1.13.0." AiCybr, November 2, 2026.

[4] Langflow. "Release v1.13.0." GitHub, October 22, 2026.

[5] Rapid7. "Metasploit Wrap Up: Langflow, MongoDB, and More." Rapid7 Blog, October 23, 2026.

[6] Swati Khandelwal. "Critical Langflow Flaws Could Let Unauthenticated Attackers Run Code on Servers." The Hacker News, November 2, 2026.

[7] CERT.UG/CC. "Langflow Unauthenticated Authentication Bypass and Remote Code Execution (CVE-2026-93680 / CVE-2026-110186)." Uganda National CERT, November 3, 2026.

[8] HOL. "Langflow CVE-2026-93680: 33 CVEs, How to Fix (Upgrade to 1.13.0)." HOL Blog, October 29, 2026.

[9] Python Package Index. "langflow: Release History." PyPI, accessed November 3, 2026.

[10] CISA. "Known Exploited Vulnerabilities Catalog." Cybersecurity and Infrastructure Security Agency, accessed November 3, 2026.

[11] Cloud Security Alliance AI Safety Initiative. "Langflow Auto-Login Bypass Chains to Unauthenticated RCE." Cloud Security Alliance, August 18, 2026.

[12] Cloud Security Alliance AI Safety Initiative. "Langflow Authorization Bypass Added to CISA's KEV Catalog." Cloud Security Alliance, July 12, 2026.

[13] Langflow. "Install Langflow." Langflow Documentation, accessed November 3, 2026.

[14] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2025.

[15] Ken Huang. "Agentic AI Threat Modeling Framework: MAESTRO." Cloud Security Alliance Blog, February 6, 2025.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.

[16] IBM. "Security Bulletin: Langflow OSS is affected by multiple vulnerabilities (1.0.0–1.12.4, fixed in 1.12.5)." IBM Support, October 19, 2026.

[17] Cloud Security Alliance AI Safety Initiative. "CVE-2026-33017: Langflow RCE Exploits Enterprise AI Pipelines." Cloud Security Alliance, 2026.

[18] Cloud Security Alliance AI Safety Initiative. "Langflow Path Traversal: Unauthenticated RCE Actively Exploited." Cloud Security Alliance, 2026.