GOVERNMENT OF THE REPUBLIC OF VANUATU

PRIME MINISTER'S OFFICE

CERTVU
DEPARTMENT OF COMMUNICATIONS & DIGITAL TRANSFORMATION

PM B 9108 Port Vila, Vanuatu
Tel: (678) 33380

GOUVERNEMENT DE LA REPUBLIQUE DU VANUATU

BUREAU DU PREMIER MINISTRE

CERTVU

DEPARTMENT DE COMMUNICATION ET DE TRANSFORMATION NUMERIQUE

SPP 9108 Port Vila, Vanuatu
Tel: (678) 33380

5 August 2026

Advisory 181: IBM Langflow Code Injection Vulnerability (CVE-2026-9198).

Release Date: 5th August 2026
Impact: CRITICAL
TLP: CLEAR

The Department of Communications and Digital Transformation (DCDT) through CERT Vanuatu (CERTVU), provides the following advisory.

This alert is relevant to Organizations and System/Network administrators that utilize the above products. This alert is intended to be understood by technical users and systems administrators.

What is it?

IBM Langflow OSS contains an Improper Control of Generation of Code vulnerability (CWE-94, CVSS v3.1 9.8 Critical, vector AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H) which allows a remote, unauthenticated attacker to achieve full remote code execution on a default Langflow deployment.

The vulnerability is a chain of two API endpoints. The /api/v1/auto_login endpoint issues SUPERUSER tokens to any network caller without authentication. The /api/v1/validate/code endpoint then executes attacker-supplied Python through exec() without sandboxing. Together these give an unauthenticated attacker arbitrary code execution as a privileged user.

Langflow is an open-source visual framework for building AI agent and retrieval-augmented generation workflows, now maintained by IBM. It is commonly deployed by development teams and innovation units, frequently as a container, and often outside formal IT asset management.

The vulnerability was published by IBM on 17 July 2026 as part of a coordinated disclosure of more than a dozen Langflow security issues. On 4 August 2026 the United States Cybersecurity and Infrastructure Security Agency (CISA) added CVE-2026-9198 to its Known Exploited Vulnerabilities (KEV) catalog, confirming active exploitation, with a federal remediation due date of 7 August 2026. CISA assessed the vulnerability as Exploitation: active, Automatable: yes, Technical Impact: total.

What are the systems affected?

The vulnerability affects IBM Langflow across the following versions:

- Langflow OSS 1.0.0 through 1.10.0 inclusive - Solution: upgrade to 1.10.1 or above

IBM Langflow, Langflow Desktop and Langflow OSS distributions are all listed as affected by the July 2026 disclosure set.

Default deployments are exploitable without any configuration change by the attacker. The AUTO_LOGIN setting, which enables the vulnerable token-issuing behaviour, is enabled by default.

Administrators should note that version 1.10.1 resolves CVE-2026-9198, but a separate issue, CVE-2026-14499, affects versions 1.0.0 through 1.10.1 inclusive. Organisations should track the current fixed release on the IBM support page rather than assuming 1.10.1 closes every Langflow issue.

What does this mean?

Background - how this vulnerability is exploited:

- An attacker locates an internet-reachable Langflow instance. No credentials, phishing or user interaction are required, and attack complexity is rated low.
- The attacker sends a request to /api/v1/auto_login, which returns a long-lived SUPERUSER bearer token to any caller when AUTO_LOGIN is enabled.
- Using that token, the attacker posts Python code to /api/v1/validate/code, which passes it to exec() without sandboxing.
- The code runs with the privileges of the Langflow process, giving full control of the host and of everything that process can reach.

In practice, the exploitation sequence is simple enough to automate at scale, which is reflected in CISA's assessment that the vulnerability is automatable and the technical impact total. Mass scanning for exposed Langflow instances should be assumed.

The consequences extend beyond the host itself. Langflow deployments routinely hold API keys and credentials for AI model providers, databases, vector stores and third-party services, stored within flows so that the platform can call them. An attacker with code execution can read all of them. Where those credentials are billable model provider keys, compromise also carries direct financial exposure.

CERTVU draws particular attention to the asset management dimension. Langflow is often installed by developers or innovation teams experimenting with AI capability, rather than by central IT, and may not appear on any organisational asset register. Organisations should not assume Langflow is absent simply because it was never formally procured.

This advisory forms part of a wider pattern that CERTVU is observing, in which AI development and orchestration platforms are being adopted rapidly and exposed to the internet without the security review normally applied to production systems.

Indicators of Compromise

There is no single definitive indicator of compromise for this issue. Administrators operating Langflow should review the following:

- Web or reverse proxy logs showing requests to /api/v1/auto_login or /api/v1/validate/code from any source outside the trusted administrative network.
- POST requests to /api/v1/validate/code containing Python code that was not submitted by a known user, particularly containing calls to os, subprocess or socket modules.
- Unexpected user accounts within Langflow, or superuser tokens issued outside normal working patterns.
- Unexpected child processes spawned by the Langflow process, including shell interpreters, and unexpected outbound network connections from the Langflow host or container.
- New cron entries, systemd services, or SSH authorized_keys additions on the Langflow host.
- Unexpected files written to the Langflow working directory or to temporary directories.
- Anomalous usage or billing activity on AI model provider accounts whose API keys are stored in Langflow flows, which may be the first visible sign of credential theft.

Mitigation process

CERTVU recommends the following:

- Upgrade Langflow to version 1.10.1 or above without delay, and redeploy any containers using the updated image. This vulnerability is on the CISA Known Exploited Vulnerabilities catalog with confirmed active exploitation and should be treated as an emergency remediation priority outside of normal patching schedules.
- Identify whether Langflow is present in your environment at all. Query asset inventories, container registries and network scans for Langflow instances, including any deployed informally by development or innovation teams.
- Remove Langflow instances from direct internet exposure. Place them behind a VPN or an authenticating reverse proxy, and restrict access to trusted source networks.
- Where immediate patching is not possible, block /api/v1/auto_login and /api/v1/validate/code at the reverse proxy for all non-trusted source networks as an interim compensating control.

- Review the Langflow configuration and disable AUTO_LOGIN, set NEW_USER_IS_ACTIVE to false, and set WEBHOOK_AUTH_ENABLE to true, as related July 2026 vulnerabilities are exploitable through these settings.
- After upgrading, verify that /api/v1/auto_login no longer issues tokens to unauthenticated callers.
- Rotate all secrets, API keys and model provider credentials stored in Langflow flows on any host that was exposed, and review those provider accounts for unauthorised usage.
- Audit affected hosts for signs of prior exploitation, including the indicators listed above. Because exploitation is confirmed and predates most patching, treat any internet-exposed instance running an affected version as potentially compromised rather than relying on the upgrade alone.
- Review the other Langflow vulnerabilities disclosed in July 2026, which include additional unauthenticated issues, path traversal, SSRF and unsafe deserialization.
- If evidence of compromise is identified, engage your internal incident response team and report the incident to CERTVU at incident@cert.gov.vu or on telephone (678) 33380.

Reference

1. https://www.cve.org/CVERecord?id=CVE-2026-9198
2. https://www.cisa.gov/known-exploited-vulnerabilities-catalog
3. https://www.cisa.gov/news-events/alerts/2026/08/04/cisa-adds-three-known-exploited-vulnerabilities-catalog
4. https://www.ibm.com/support/pages/node/7278927
5. https://nvd.nist.gov/vuln/detail/CVE-2026-9198
6. https://cwe.mitre.org/data/definitions/94.html