Hive Pro

Threat Level

Red

Hiveforce Labs

THREAT ADVISORY

VULNERABILITY REPORT

From Disclosure to Exploitation Overnight of a CVE-2026-33017 Langflow Flaw

Date of Publication	Admiralty Code	TA Number
March 23, 2026	A1	TA2026080

Summary

First Seen: March 17, 2026
Affected Product: Langflow
Impact: CVE-2026-33017 is a critical unauthenticated remote code execution (RCE) vulnerability affecting Langflow, the popular open-source visual framework for building AI agents and Retrieval-Augmented Generation (RAG) pipelines. The impact of CVE-2026-33017 is severe and far-reaching. Successful exploitation grants an attacker full server process privileges, enabling arbitrary command execution and complete system compromise via a single unauthenticated HTTP POST request. The extremely low barrier to exploitation, no authentication, no multi-step chains, and a simple JSON payload, combined with the large attack surface of publicly exposed Langflow instances, make this vulnerability an urgent priority for immediate remediation.

CVE

CVE	NAME	AFFECTED PRODUCT	ZERO-DAY	CISA KEV	PATCH
CVE-2026-33017	Langflow Unauthenticated Remote Code Execution Vulnerability	Langflow	✗	✗	✓

Vulnerability Details

#1

A critical vulnerability was discovered in Langflow, an open-source visual framework used to build AI agents and Retrieval-Augmented Generation (RAG) pipelines. Within 20 hours of public disclosure, attackers had already begun exploiting the flaw, demonstrating how quickly newly revealed security issues are turned into real-world attacks.

#2

The vulnerability, tracked as CVE-2026-33017, is an unauthenticated remote code execution flaw in Langflow's public flow build endpoint. It allows an attacker to execute arbitrary Python code on any exposed instance without needing credentials. The attack requires only a single HTTP request, making exploitation fast and simple.

#3

The issue affects the endpoint 'POST /api/v1/build_public_tmp/{flow_id}/flowendpoint', which is intended to let users build public flows without authentication. The endpoint accepts user-supplied flow data that may contain Python code inside node definitions. Because this code is executed server-side without proper isolation or sandboxing, an attacker can run malicious code directly on the host system.

#4

This design makes the vulnerability especially dangerous. The endpoint is publicly accessible by default, so automated scanning and mass exploitation are trivial. Attackers quickly developed working exploits and began scanning the internet for exposed Langflow instances shortly after disclosure.

#5

Compromised systems have already shown signs of data theft, including the extraction of API keys, credentials, and other sensitive information. Since Langflow is often configured with access to services such as cloud platforms, language model APIs, and databases, a single breach can provide attackers with broader access to infrastructure and data, increasing the risk of supply chain compromise.

#6

The scale of the risk is amplified by Langflow's popularity. With over 145,000 stars on GitHub, the framework has a wide user base, which translates into a large number of potentially exposed deployments. This significantly expands the available attack surface. CVE-2026-33017 reflects a growing trend in cybersecurity, critical vulnerabilities in widely used open-source tools are now weaponized within hours of disclosure. Attackers no longer wait for public proof-of-concept code. Instead, they analyze patches and advisories immediately, develop their own exploits, and begin targeting systems almost in real time.

Vulnerability

CVE ID	AFFECTED PRODUCTS	AFFECTED CPE	CWE ID
CVE-2026-33017	Langflow Langflow version before 1.9.0	cpe:2.3:a:langflow:langflow:*:*:*:*:*:*:*:*	CWE-94, CWE-95, CWE-306

Recommendations

Update Langflow Immediately: Upgrade all Langflow instances to version 1.9.0 or later without delay. This patched version removes the ability for the unauthenticated endpoint to accept attacker-supplied flow data containing arbitrary executable code. Given that active exploitation was observed within 20 hours of disclosure, this update is of the highest urgency.

Restrict Network Access to Langflow Instances: Langflow should not be directly exposed to the internet without an authentication layer. Implement firewall rules or deploy a reverse proxy with authentication in front of all Langflow instances. Specifically, restrict access to the `/api/v1/build_public_tmp` endpoint or disable public flow building entirely if not required.

Audit and Rotate Credentials: Immediately audit environment variables, API keys, database passwords, and cloud credentials on any publicly exposed Langflow instance. Rotate all secrets as a precaution, as observed attackers specifically targeted environment variable dumps and `.env` file extraction to harvest OpenAI, Anthropic, AWS, and database credentials.

Disable AUTO_LOGIN in Production: The default `AUTO_LOGIN=true` configuration allows unauthenticated users to obtain superuser tokens, which dramatically lowers the exploitation barrier. Disable this setting in any production or internet-facing deployment and enforce proper authentication controls.

Monitor for Post-Exploitation Indicators: Monitor for outbound connections to unusual ports or known callback services such as `oastify.com`, `interact.sh`, `oast.live`, `oast.me`, `oast.pro`, and `dnslog.cn`, which indicate active exploitation and data exfiltration. Additionally, monitor for unexpected process execution (shell commands spawned from the Langflow process), reads of sensitive files like `/etc/passwd` or `.env`, and outbound HTTP connections to unfamiliar IP addresses.

Potential MITRE ATT&CK TTPs

Tactic	Technique	Sub-technique
Initial Access	T1190: Exploit Public-Facing Application	
Execution	T1059: Command and Scripting Interpreter	T1059.006: Python
		T1059.004: Unix Shell
Discovery	T1082: System Information Discovery	
	T1083: File and Directory Discovery	
Credential Access	T1552: Unsecured Credentials	T1552.001: Credentials In Files

Tactic	Technique	Sub-technique
Collection	T1005: Data from Local System	
Exfiltration	T1041: Exfiltration Over C2 Channel	
Command and Control	T1071: Application Layer Protocol	T1071.001: Web Protocols
	T1105: Ingress Tool Transfer	

Indicators of Compromise (IOCs)

TYPE	VALUE
IPv4	77[.]110[.]106[.]154, 209[.]97[.]165[.]247, 188[.]166[.]209[.]86, 205[.]237[.]106[.]117, 83[.]98[.]164[.]238, 173[.]212[.]205[.]251
IPv4:Port	143[.]110[.]183[.]86[:]8080, 173[.]212[.]205[.]251[:]8443
URL	hxxp[:]//143[.]110[.]183[.]86[:]8080/, hxxp[:]//173[.]212[.]205[.]251[:]8443/z

Patch Link

https://github.com/langflow-ai/langflow/releases

References

https://www.sysdig.com/blog/cve-2026-33017-how-attackers-compromised-langflow-ai-pipelines-in-20-hours

https://github.com/langflow-ai/langflow/security/advisories/GHSA-vwmf-pq79-vjvx

https://github.com/advisories/GHSA-rvqx-wpfh-mfx7

https://github.com/langflow-ai/langflow/commit/73b6612e3ef25fdae0a752d75b0fabd47328d4f0

What Next?

At **Hive Pro**, it is our mission to detect the most likely threats to your organization and to help you prevent them from happening.

Book a free demo with **HivePro Uni5**: Threat Exposure Management Platform.

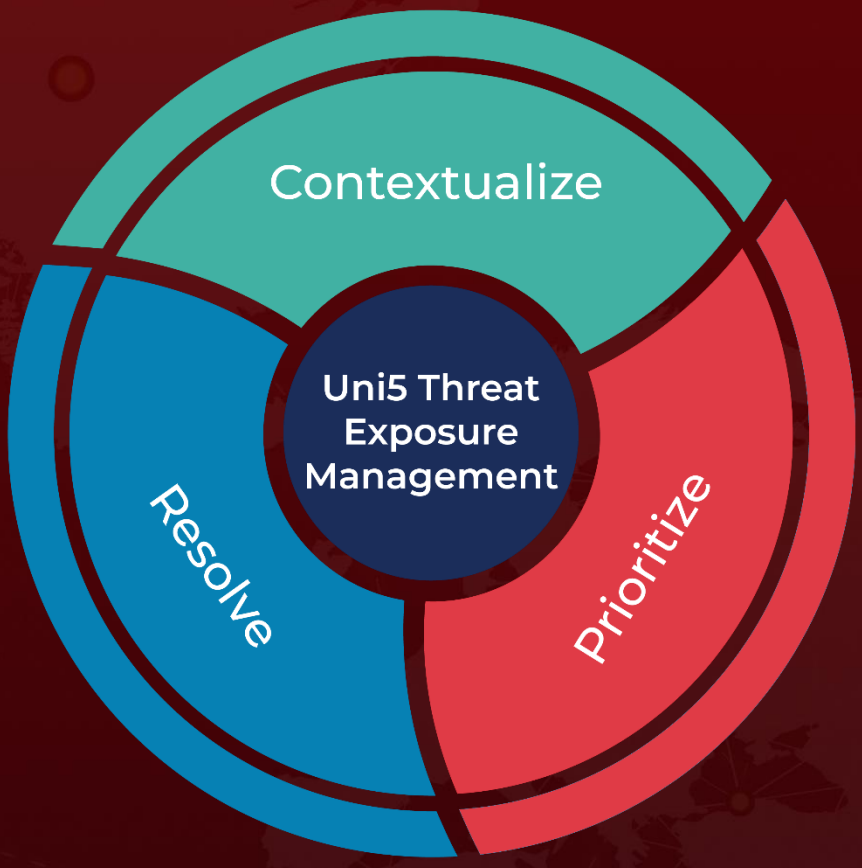

REPORT GENERATED ON

March 23, 2026 • 08:30 AM

© 2026 All Rights are Reserved by Hive Pro

More at www.hivepro.com