INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-07-06 06:34 UTC

# IBM Langflow OSS Secret-Reading Vulnerabilities (CVE-2026-10134)

CVE VULNERABILITY | CRITICAL

| | |
|---|---|
| SCC Item ID | SCC-CVE-2026-0389 |
| Type | CVE Vulnerability |
| CVE ID | CVE-2026-10134 |
| Severity | CRITICAL |
| EPSS Score | 0.0031 (23th percentile) |
| Affected Products | IBM Langflow OSS (specific version range not confirmed from available sources) |
| Published | 2026-07-06 |
| Discovery Source | Gemini |

## Executive Summary

IBM Langflow OSS is reported to be affected by a vulnerability tracked as CVE-2026-10134, which, according to a single aggregated news source, could allow an attacker to read active process secrets and modify AI workflows without authorization. Affected version ranges have not been confirmed by an IBM security advisory or verified NVD enrichment at the time of writing. Until IBM publishes a first-party advisory, organizations running IBM Langflow OSS should treat this as an unconfirmed but credible risk requiring immediate inventory and monitoring.

## Technical Analysis

CVE-2026-10134 is reported to affect IBM Langflow OSS, an open-source AI workflow orchestration platform. According to the single available source (dailycybersecurity.com, a T3 aggregator), the vulnerability allegedly enables an attacker to read every active process secret and modify AI workflows without restriction. MITRE ATT&CK techniques T1552 (Unsecured Credentials) and T1565 (Data Manipulation) are associated with the reported behavior. No CVSS base score, CVSS vector, or CWE assignment is available in the provided data; the qualitative rating of 'critical' originates from the source item metadata, not from a verified NVD or IBM advisory. EPSS score is 0.00314 (23rd percentile), indicating low current exploitation probability based on available signals. CISA KEV inclusion is not confirmed. Affected version range is unconfirmed. No IBM Product Security Incident Response Team (PSIRT) advisory or NVD detail enrichment was present in the source material. Confidence in all technical specifics is LOW pending first-party confirmation.

## Action Checklist

1. Step 1: Containment, Inventory all instances of IBM Langflow OSS running in your environment immediately. If the service is internet-facing, restrict inbound access to known trusted IP ranges or place it behind an authenticated reverse proxy until IBM publishes a confirmed advisory and patch. Reference NIST AC-17 (Remote Access) for connection restriction guidance.
2. Step 2: Detection, Query application and process logs on Langflow OSS hosts for anomalous reads of environment variables, secrets files, or credential stores, and for unexpected modifications to workflow definitions. Monitor for MITRE T1552-related behavior (access to /proc, environment variable enumeration, secrets manager API calls) and T1565-related behavior (unauthorized workflow configuration changes). Reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) and CIS 8.2 (Collect Audit Logs).
3. Step 3: Eradication, Monitor IBM PSIRT (https://www.ibm.com/support/pages/ibm-security-advisories) for a confirmed advisory and patch. No vendor-confirmed patch version or remediation path is available in the provided source material; do not apply speculative fixes. Once a patch is issued, follow IBM's documented upgrade path. Reference CIS 7.3 (Perform Automated Operating System Patch Management) and CIS 7.4 (Perform Automated Application Patch Management).
4. Step 4: Recovery, After applying any IBM-issued patch, rotate all secrets, API keys, and credentials accessible to Langflow OSS processes (D3-CRO: Credential Rotation). Audit AI workflow definitions for unauthorized modifications introduced during the exposure window. Re-enable full production access only after secrets rotation is confirmed complete and workflow integrity is verified. Reference NIST AC-6 (Least Privilege) to scope secret access going forward.
5. Step 5: Post-Incident, Review how Langflow OSS is granted access to secrets and whether least-privilege principles (NIST AC-6) are enforced on AI workflow service accounts. Evaluate whether secrets are stored in a dedicated vault with audit logging rather than as process environment variables. Conduct a post-incident review against CIS 7.1 (Establish and Maintain a Vulnerability Management Process) to ensure AI/ML platform components are included in your vulnerability management scope, as these are frequently excluded from standard asset inventories.

## IR / Forensic Enrichment

| | |
|---|---|
| **Triage Priority** | URGENT |
| **Escalation Criteria** | Escalate to CISO and legal/privacy counsel immediately if Langflow application logs, auditd records, or secrets manager audit trails show evidence of successful secret exfiltration or unauthorized workflow modification, as exposed credentials may govern downstream AI pipelines processing PII, PHI, or regulated data, triggering breach notification obligations under applicable privacy regulations. |
| **Recovery Notes** | After patching and secrets rotation, monitor Langflow API access logs and integrated secrets manager audit trails for a minimum of 30 days for re-use of any rotated credentials or re-emergence of unauthorized workflow modifications, as an attacker who exfiltrated secrets before containment may attempt to re-establish access via a different vector. Verify workflow integrity by comparing all flow definitions against the last known-good backup and confirm that no externally-controlled model endpoints, webhook URLs, or API key references were injected. Do not restore internet-facing access until IBM confirms the patched version resolves CVE-2026-10134 and internal smoke testing of the upgraded instance is complete. |

| | |
|---|---|
| **Forensic Artifacts** | Langflow process environment dump (/proc//environ on Linux): contains the live secrets — API keys, database credentials, model provider tokens — that CVE-2026-10134 reportedly exposes to an attacker; must be captured before any process termination or host isolation \| Langflow workflow database (langflow.db, SQLite, at the path configured in LANGFLOW_DATABASE_URL): records all workflow definitions including model endpoints, tool configurations, and API key references; unauthorized POST/PUT to /api/v1/flows/ would persist here and is the primary artifact for detecting T1565-style workflow tampering \| Langflow application log (langflow.log in the working directory or as configured): contains timestamped HTTP request records including source IPs, endpoints accessed, and HTTP methods; pivot on PUT/POST to /api/v1/flows/ and any GET requests to secret-retrieval endpoints during the exposure window \| auditd logs or Sysmon Event ID 10/11 records on the Langflow host: capture file open syscalls against /proc/*/environ (Linux) and FileCreate/ProcessAccess events (Windows) that indicate active secret enumeration consistent with the reported vulnerability mechanism \| Integrated secrets manager audit trail (AWS CloudTrail GetSecretValue events, HashiCorp Vault audit log, or equivalent): records whether secrets accessible to the Langflow process were actually retrieved via API call from an external or unexpected source IP during the exposure window, distinguishing confirmed exfiltration from theoretical exposure |

### Per-Action IR Details

**Step 1: Containment — Inventory all instances of IBM Langflow OSS running in your environment immediately. If the service is internet-facing, restrict inbound access to known trusted IP ranges or place it behind an authenticated reverse proxy until IBM publishes a confirmed advisory and patch. Reference NIST AC-17 (Remote Access) for connection restriction guidance.**

**NIST Phase:** Containment

**Reference:** NIST 800-61r3 §3.3 — Containment Strategy

**Controls:** NIST AC-17 (Remote Access), NIST AC-4 (Information Flow Enforcement), CIS 4.4 (Implement and Manage a Firewall on Servers), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory)

**Compensating:** Run 'ss -tlnp | grep ' or 'netstat -ano | findstr ' on each host to identify listening Langflow OSS processes. Use iptables (Linux) or Windows Firewall with Advanced Security to restrict inbound access to the Langflow API port (default 7860) to a known-good IP allowlist. Document all discovered instances in a spreadsheet before restricting access. A 2-person team can complete a scan across a /24 subnet using 'nmap -p 7860 --open 192.168.1.0/24' in under 10 minutes.

**Evidence:** Before isolating or restricting any Langflow OSS instance, capture: (1) active network connections from the Langflow process ('ss -antp' or 'Get-NetTCPConnection' filtered to the Langflow PID), preserving any attacker C2 or exfiltration endpoints; (2) a full process listing with parent-child relationships ('ps auxf' on Linux or 'Get-Process | Select-Object Id,Name,Path,StartTime' on Windows) to identify anomalous child processes spawned by Langflow; (3) current environment variables of the Langflow process ('/proc//environ' on Linux, or use Sysinternals Process Explorer to dump the process environment) — these contain the secrets CVE-2026-10134 reportedly exposes and will be destroyed upon process kill or host isolation.

**Step 2: Detection — Query application and process logs on Langflow OSS hosts for anomalous reads of environment variables, secrets files, or credential stores, and for unexpected modifications to workflow definitions. Monitor for MITRE T1552-related behavior (access to /proc, environment variable enumeration, secrets manager API calls) and T1565-related behavior (unauthorized workflow configuration changes). Reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) and CIS 8.2 (Collect Audit Logs).**

**NIST Phase:** Detection Analysis

**Reference:** NIST 800-61r3 §3.2 — Detection and Analysis

**Controls:** NIST AU-6 (Audit Record Review, Analysis, And Reporting), NIST AU-2 (Event Logging), NIST AU-12 (Audit Record Generation), CIS 8.2 (Collect Audit Logs)

**Compensating:** On Linux hosts: enable auditd with rules targeting '/proc/*/environ' reads ('auditctl -a always,exit -F arch=b64 -S open,openat -F path=/proc -k langflow_secret_read') and Langflow's workflow storage directory (typically '~/.langflow/' or the configured LANGFLOW_CONFIG_DIR). On Windows: deploy Sysmon with a config that captures Event ID 10 (ProcessAccess) targeting the Langflow process and Event ID 11 (FileCreate) in the workflow config path. Query Langflow's own application log (default: 'langflow.log' in the working directory) for HTTP POST/PUT requests to '/api/v1/flows/' endpoints, which indicate workflow modification attempts. Pipe results through 'grep -E "(PUT|POST).*/flows/"' for a quick triage pass.

**Evidence:** This is a detection step that does not alter live state — no volatile pre-capture is strictly required before querying logs. However, if the Langflow process is still running during analysis, preserve '/proc//environ' and '/proc//fd/' (open file descriptors) immediately, as they provide ground truth on which secrets the process currently holds and which files it has open. Also snapshot Langflow's workflow database file (SQLite at the path set in LANGFLOW_DATABASE_URL, default 'langflow.db') before any write activity occurs, to establish a baseline for identifying unauthorized workflow modifications introduced during the exposure window.

**Step 3: Eradication — Monitor IBM PSIRT (https://www.ibm.com/support/pages/ibm-security-advisories) for a confirmed advisory and patch. No vendor-confirmed patch version or remediation path is available in the provided source material; do not apply speculative fixes. Once a patch is issued, follow IBM's documented upgrade path. Reference CIS 7.3 (Perform Automated Operating System Patch Management) and CIS 7.4 (Perform Automated Application Patch Management).**

**NIST Phase:** Eradication

**Reference:** NIST 800-61r3 §3.4 — Eradication

**Controls:** CIS 7.3 (Perform Automated Operating System Patch Management), CIS 7.4 (Perform Automated Application Patch Management), CIS 7.2 (Establish and Maintain a Remediation Process), CIS 2.2 (Ensure Authorized Software is Currently Supported)

**Compensating:** Until IBM PSIRT publishes a confirmed advisory, subscribe to the IBM Security Advisories RSS feed (https://www.ibm.com/support/pages/ibm-security-advisories) and set a daily calendar reminder to manually check for CVE-2026-10134. As a compensating measure, restrict the Langflow OSS process account to a dedicated low-privilege service account with no access to secret stores beyond what Langflow strictly requires ('useradd -r -s /sbin/nologin langflow'). Do not upgrade Langflow speculatively to an unverified version — confirm the fixed version from the IBM advisory before applying. Document the patch hold decision with a dated risk acceptance note signed by the system owner.

**Evidence:** Before applying any IBM-issued patch or performing an upgrade, capture: (1) a full memory dump of the running Langflow process ('gcore ' on Linux or using ProcDump: 'procdump.exe -ma ') to preserve in-memory secrets and any attacker-injected workflow logic that may not persist to disk; (2) a copy of the current Langflow workflow database ('langflow.db') and any exported flow JSON files in LANGFLOW_CONFIG_DIR to document the pre-patch workflow state; (3) the installed package version ('pip show langflow' or 'pip3 show langflow') and all dependency versions ('pip freeze > langflow_preupgrade_deps.txt') to support post-patch regression analysis and vendor confirmation of the remediated version.

**Step 4: Recovery — After applying any IBM-issued patch, rotate all secrets, API keys, and credentials accessible to Langflow OSS processes (D3-CRO: Credential Rotation). Audit AI workflow definitions for unauthorized modifications introduced during the exposure window. Re-enable full production access only after secrets rotation is confirmed complete and workflow integrity is verified. Reference NIST AC-6 (Least Privilege) to scope secret access going forward.**

**NIST Phase:** Recovery

**Reference:** NIST 800-61r3 §3.5 — Recovery

**Controls:** NIST AC-6 (Least Privilege), NIST AC-2 (Account Management), CIS 6.2 (Establish an Access Revoking Process), CIS 5.2 (Use Unique Passwords)

**Compensating:** Enumerate all secrets accessible to the Langflow process by reviewing: (1) the '.env' file or environment variable injection in the Langflow systemd unit or Docker Compose file; (2) any secrets manager integrations configured in Langflow (e.g., AWS Secrets Manager, HashiCorp Vault) by inspecting Langflow's config YAML or component configuration. Rotate each credential individually, confirming revocation of old values before re-injection. For workflow integrity verification, export all flow definitions as JSON ('GET /api/v1/flows/' against the now-patched instance) and diff them against a known-good backup using 'diff -u backup_flows.json current_flows.json' — focus on changes to model endpoint URLs, API key references, and tool/action nodes, which are the fields an attacker exploiting T1565-style workflow tampering would most likely modify.

**Evidence:** Before rotating credentials or re-enabling production access, preserve: (1) the complete set of current workflow definitions exported from the Langflow API as timestamped JSON files — these document what was accessible and potentially tampered during the exposure window; (2) all Langflow application logs covering the exposure period (from initial deployment or last known-good state to containment timestamp), archived to write-once storage to support later forensic review; (3) audit logs from any integrated secrets managers or cloud credential stores showing API calls originating from the Langflow host's IP during the exposure window, to determine whether exposed secrets were actually exfiltrated or used externally.

**Step 5: Post-Incident — Review how Langflow OSS is granted access to secrets and whether least-privilege principles (NIST AC-6) are enforced on AI workflow service accounts. Evaluate whether secrets are stored in a dedicated vault with audit logging rather than as process environment variables. Conduct a post-incident review against CIS 7.1 (Establish and Maintain a Vulnerability Management Process) to ensure AI/ML platform components are included in your vulnerability management scope, as these are frequently excluded from standard asset inventories.**

**NIST Phase:** Post Incident

**Reference:** NIST 800-61r3 §4 — Post-Incident Activity

**Controls:** NIST AC-6 (Least Privilege), NIST AU-9 (Protection Of Audit Information), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 2.1 (Establish and Maintain a Software Inventory)

**Compensating:** Conduct a 60-minute structured lessons-learned session (NIST 800-61r3 §4 recommends within 2 weeks of resolution) specifically addressing: (1) whether IBM Langflow OSS appeared in your asset inventory and vulnerability scan scope prior to this event — if not, add it and all other AI/ML framework deployments (LangChain, Ollama, LocalAI, etc.) to your asset register and scan policy; (2) whether secrets injected into Langflow processes are sourced from environment variables in unit files or Compose files (high risk) versus a vault solution with per-secret audit logging (preferred). As a free vault option for small teams, evaluate HashiCorp Vault Community Edition or 'pass' (the Unix password manager) with GPG encryption as a step up from plaintext env files. Document findings and owners in a remediation tracking ticket with a 30-day deadline.

**Evidence:** For the post-incident record, assemble: (1) a timeline reconstructed from Langflow application logs, auditd/Sysmon logs, and secrets manager audit trails covering the full exposure window; (2) the diffed workflow definitions from Step 4 showing any confirmed unauthorized modifications; (3) evidence of which specific secrets were in scope (from the pre-rotation env file review) versus which were confirmed accessed by an external party (from secrets manager audit logs) — this distinction is material if breach notification obligations apply. Retain all artifacts per your documented retention policy (NIST AU-11) and store in write-protected, access-controlled evidence storage.

## Detection Guidance

No confirmed IOCs, CVE-specific signatures, or vendor-issued detection rules are available in the provided source material. The following guidance is based on the reported behavior and associated MITRE techniques. On Langflow OSS hosts, monitor process logs for: (1) T1552 indicators, bulk reads of environment variables, access to secrets files or credential stores (e.g., /proc/[pid]/environ on Linux), or unusual calls to secrets management APIs; (2) T1565 indicators, modifications to workflow definition files or database records outside of

normal deployment windows, especially from network-sourced sessions. Enable audit logging per NIST AU-2 (Event Logging) and AU-12 (Audit Record Generation) on all Langflow OSS hosts if not already active. Cross-reference D3-SFA (System File Analysis) to monitor workflow configuration files for unauthorized changes. Note: no confirmed exploit signatures, network indicators, or affected endpoint patterns can be provided from the available sources. Detection rules should be treated as hypothesis-driven hunting until IBM releases a confirmed advisory.

## Framework Mappings

### MITRE-ATTACK

- **T1552** — Unsecured Credentials
- **T1565** — Data Manipulation

### ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities

### NIST-CSF-2

- **DE.AE-08** — Incidents are declared when adverse events meet the defined incident criteria

### NIST-800-53R5

- **IR-5** — Incident Monitoring

## MITRE ATT&CK Mapping

| Technique ID | Technique Name | Tactic |
|---|---|---|
| T1552 | Unsecured Credentials | Credential-Access |
| T1565 | Data Manipulation | Impact |

## Sources

| Source | URL | Tier |
|---|---|---|
| **gemini** | https://dailycybersecurity.com/vulnerability-intelligence-report-ju... | T3 |
| **CVE-2026-10134 Detail - NVD** | https://nvd.nist.gov/vuln/detail/CVE-2026-10134 | T1 |
| **CVE-2026-10134 - Exploits & Severity - Feedly** | https://feedly.com/cve/CVE-2026-10134 | T3 |
| **Vulnerability database \| ManageEngine Vulnerability Manager Plus** | https://www.manageengine.com/vulnerability-management/vulnerability... | T3 |

| Source | URL | Tier |
| --- | --- | --- |
| **Known Exploited Vulnerabilities Catalog \| CISA** | https://www.cisa.gov/known-exploited-vulnerabilities-catalog | **T1** |



Generated 2026-07-06 06:34 UTC by TJS Security Command Center