INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-26 18:38 UTC

Langflow IDOR Vulnerability Allows Authenticated Attackers to Execute Arbitrary User Flows (CVE-2026-55255)

CVE VULNERABILITY | **HIGH** | CVSS 8.5 | **CISA KEV**

SCC Item ID	SCC-CVE-2026-0353
Type	CVE Vulnerability
CVE ID	CVE-2026-55255
Severity	HIGH
CVSS Base Score	8.5
EPSS Score	0.0023 (14th percentile)
KEV Status	Yes — CISA Known Exploited Vulnerability
Affected Products	langflow/langflow < 1.9.2
Published	2026-06-26T00:00:00Z
Discovery Source	Vulncheck Kev

Executive Summary

A high-severity access control flaw in Langflow, an open-source platform used to build and deploy AI agent workflows, allows authenticated users to execute AI flows belonging to other users without authorization. Any organization running Langflow versions prior to 1.9.2 is exposed, and both CISA and VulnCheck confirm active exploitation in the wild. Immediate patching is required; unpatched instances risk unauthorized manipulation of AI-driven business processes and potential data exposure across tenant boundaries.

Technical Analysis

CVE-2026-55255 is an Insecure Direct Object Reference (IDOR) vulnerability (CWE-639) in the /api/v1/responses endpoint of Langflow (langflow/langflow < 1.9.2). The endpoint accepts a flow ID parameter but performs no authorization check to verify the requesting user owns or holds permission to execute the referenced flow. An authenticated attacker supplies an arbitrary victim flow ID in a crafted API request to trigger execution of that flow. CVSS base score: 8.5 (High). EPSS: 0.233% (14th percentile). MITRE techniques: T1565.001 (Stored Data Manipulation) and T1078 (Valid Accounts). Active exploitation confirmed by CISA KEV and VulnCheck KEV. Remediation: upgrade to Langflow 1.9.2 or later.

Action Checklist

1. Step 1: Containment, Immediately restrict external access to the /api/v1/responses endpoint via WAF or network ACL rules for all Langflow instances running versions prior to 1.9.2. If internet-facing, take the service offline or place it behind a VPN until patching is complete. Per CIS Controls v8 4.4, enforce host-based firewall rules to block unauthorized inbound connections to the Langflow service port.
2. Step 2: Detection, Query API gateway and application logs for requests to /api/v1/responses where the authenticated user's ID does not match the owner of the submitted flow ID. Look for high-frequency or cross-user flow execution patterns. Review audit logs (NIST AU-6) for anomalous API call volumes from single accounts, especially calls referencing flow IDs not created by that account. Flag any flow execution events occurring outside of normal business hours or from unfamiliar source IPs.
3. Step 3: Eradication, Upgrade all Langflow deployments to version 1.9.2 or later, which introduces authorization checks on the /api/v1/responses endpoint. Verify the upgrade by confirming the version string in the Langflow admin interface or via 'pip show langflow'. After patching, rotate API keys and session tokens for all Langflow service accounts per NIST AC-2 (Account Management) to invalidate any sessions established during the exposure window.
4. Step 4: Recovery, After patching, re-enable access to the endpoint and monitor /api/v1/responses traffic for continued anomalous cross-user flow execution attempts. Validate that authorization checks are enforced by testing with a non-owner account attempting to execute another user's flow ID, the request should return a 403. Enable enhanced logging per NIST AU-2 and AU-12 on the endpoint for a minimum of 30 days post-remediation. Review all flow execution logs from the exposure period to identify any flows that were executed by unauthorized users.
5. Step 5: Post-Incident, Conduct a review of all API endpoints in Langflow and any adjacent AI workflow platforms for similar missing object-level authorization checks. Map findings against NIST IR-4 (Incident Handling) to update the incident response playbook with IDOR-specific detection and containment steps. Evaluate whether AI workflow platforms in your environment are subject to periodic authorization audits. Reference NIST AC-3 (Access Enforcement) and CIS Controls v8 6.2 to verify that access granting and revoking processes account for AI agent and workflow platform permissions, not just traditional IT systems.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to legal, privacy, and executive leadership immediately if log analysis from Step 2 confirms that any unauthorized cross-user flow execution accessed flows containing PII, PHI, API credentials, or proprietary AI model configurations, as this may trigger breach notification obligations under GDPR, HIPAA, or applicable state privacy laws; also escalate if the organization cannot patch or isolate within 4 hours given CISA-confirmed active exploitation in the wild.

Recovery Notes	After deploying Langflow 1.9.2, validate authorization enforcement on /api/v1/responses with explicit cross-user token tests before restoring external access, and confirm the Langflow admin interface reflects version 1.9.2 or later. Monitor API gateway and Langflow application logs daily for 30 days post-remediation for any 200-response POST requests to /api/v1/responses involving flow IDs not owned by the authenticated caller, which would indicate either incomplete patching or a novel bypass. Review the full flow execution history from the exposure window to determine which AI workflows were executed without authorization, assess what data or downstream systems those workflows touched, and document findings for the post-incident report and any required regulatory notifications.
Forensic Artifacts	Langflow application logs (/path/to/langflow.log or container stdout): POST /api/v1/responses entries with flow_id parameter values — cross-reference each against the flow table's user_id column to identify unauthorized cross-user execution events specific to this IDOR flaw. | Langflow database flow_run table (SQLite or PostgreSQL): records the executing user_id, target flow_id, execution timestamp, and output for every flow invocation — the primary forensic record of which flows were triggered by non-owner accounts during the CVE-2026-55255 exposure window. | API gateway access logs (nginx access.log, AWS ALB logs, or equivalent): source IP addresses, user-agent strings, HTTP response codes, and request timestamps for /api/v1/responses — enables attribution of exploitation attempts to specific external clients or threat actors. | Langflow database api_key and session tables: API keys and session tokens active during the exposure window, with creation timestamps and associated user accounts — required to scope credential rotation and identify accounts whose tokens may have been harvested or reused. | Network flow records (NetFlow, VPC Flow Logs, or `ss`/`netstat` output captured at isolation): external IP addresses connected to the Langflow service port at time of containment — correlates with API gateway logs to identify scanning or automated exploitation activity targeting the /api/v1/responses endpoint.

Per-Action IR Details

Step 1: Containment — Immediately restrict external access to the /api/v1/responses endpoint via WAF or network ACL rules for all Langflow instances running versions prior to 1.9.2. If internet-facing, take the service offline or place it behind a VPN until patching is complete. Per CIS 4.4, enforce host-based firewall rules to block unauthorized inbound connections to the Langflow service port.

NIST Phase: Containment

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

Controls: CIS 4.4 (IG1/IG2/IG3) — Implement and Manage a Firewall on Servers, CIS 4.5 (IG1/IG2/IG3) — Implement and Manage a Firewall on End-User Devices

Compensating: On Linux hosts, run `iptables -I INPUT -p tcp --dport 7860 -j DROP` (or the configured Langflow port) to immediately block inbound connections to the Langflow service. For WAF-less environments, use nginx or Caddy with an IP allowlist restricting /api/v1/responses to internal CIDR ranges only. Document the block timestamp for the exposure window calculation.

Evidence: Before isolating or restricting the service, capture: (1) active TCP connections to the Langflow port via `ss -tnp sport = :7860` or `netstat -ano | findstr :7860` to record connected client IPs at time of isolation; (2) running Langflow process details via `ps aux | grep langflow` or `Get-Process` on Windows to confirm version and process tree; (3) current Langflow application log tail (`tail -n 500 /path/to/langflow.log`) to preserve the most recent /api/v1/responses request activity before log rotation or service restart truncates it. Volatile network state is destroyed the moment the service is firewalled or taken offline.

Step 2: Detection — Query API gateway and application logs for requests to /api/v1/responses where the authenticated user's ID does not match the owner of the submitted flow ID. Look for high-frequency or cross-user flow execution patterns. Review audit logs (NIST AU-6) for anomalous API call volumes from

single accounts, especially calls referencing flow IDs not created by that account. Flag any flow execution events occurring outside of normal business hours or from unfamiliar source IPs.

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-3 — Content Of Audit Records

Compensating: Without a SIEM, parse Langflow application logs directly using: `grep 'POST /api/v1/responses' /path/to/langflow.log | awk '{print $1, $7, $9}' | sort | uniq -c | sort -rn` to surface high-frequency callers. Cross-reference the `flow_id` parameter in each POST body against Langflow's SQLite or PostgreSQL `flow` table (`SELECT id, user_id FROM flow`) to identify requests where the calling user's ID does not match the flow owner's `user_id`. Use `jq` to parse JSON-formatted logs: `cat langflow.log | jq 'select(.path=="/api/v1/responses") | {user: .user_id, flow: .body.flow_id, ip: .remote_addr}'`.

Evidence: This step is analytical and does not alter live state, so no volatile pre-capture is required before querying logs. However, preserve log files to read-only or copy them to evidence storage before analysis to prevent accidental modification: `cp -p /path/to/langflow.log /evidence/langflow_$(date +%Y%m%d%H%M%S).log`. Key artifacts to examine: Langflow application logs containing POST /api/v1/responses entries with mismatched user-to-flow-owner mappings; API gateway access logs (e.g., nginx access.log, AWS ALB access logs) showing source IPs, user-agent strings, and response codes for /api/v1/responses; Langflow database flow execution history table recording which user_id triggered each flow_id and the resulting output.

Step 3: Eradication — Upgrade all Langflow deployments to version 1.9.2 or later, which introduces authorization checks on the /api/v1/responses endpoint. Verify the upgrade by confirming the version string in the Langflow admin interface or via 'pip show langflow'. After patching, rotate API keys and session tokens for all Langflow service accounts per D3-CRO (Credential Rotation) to invalidate any sessions established during the exposure window.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST SI-1 — Policy And Procedures, CIS 7.3 (IG1/IG2/IG3) — Perform Automated Operating System Patch Management, CIS 7.4 (IG1/IG2/IG3) — Perform Automated Application Patch Management

Compensating: Run `pip install --upgrade langflow==1.9.2` in the Langflow virtual environment, then validate with `pip show langflow | grep Version`. For containerized deployments, pull the updated image (`docker pull langflowai/langflow:1.9.2`) and redeploy, confirming the version endpoint at `GET /health` or `/version`. To rotate API keys without an identity platform, query the Langflow database: `UPDATE api_key SET key = gen_random_uuid() WHERE user_id IN (SELECT id FROM user)` and force re-login by invalidating all active session tokens in the sessions table.

Evidence: Before applying the patch or rotating credentials, capture: (1) a full export of the Langflow database's `flow`, `user`, `api_key`, and `flow_run` tables to preserve the pre-patch state showing all flows executed during the exposure window and their triggering user accounts — this is the primary forensic record of unauthorized cross-user flow executions; (2) active session tokens currently stored in the database sessions table (`SELECT * FROM session`) before rotation destroys evidence of sessions established by exploiting the IDOR flaw; (3) the installed package manifest (`pip freeze > /evidence/pip_freeze_pre_patch.txt`) to document the vulnerable version at time of remediation. Credential rotation permanently invalidates session evidence — capture it first.

Step 4: Recovery — After patching, re-enable access to the endpoint and monitor /api/v1/responses traffic for continued anomalous cross-user flow execution attempts. Validate that authorization checks are enforced by testing with a non-owner account attempting to execute another user's flow ID — the request should return a 403. Enable enhanced logging per NIST AU-2 and AU-12 on the endpoint for a minimum of 30 days post-remediation. Review all flow execution logs from the exposure period to identify any flows that were executed by unauthorized users.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST AU-2 — Event Logging, NIST AU-12 — Audit Record Generation, NIST AU-6 — Audit Record Review, Analysis, And Reporting, NIST AU-4 — Audit Storage Capacity

Compensating: Perform authorization validation using curl: `curl -X POST https:///api/v1/responses -H 'Authorization: Bearer ' -d '{"flow_id": ""}'` and confirm the response is HTTP 403. Enable verbose request logging in Langflow by setting `LANGFLOW_LOG_LEVEL=DEBUG` in the environment, directing output to a retained log file. Set up a cron job to alert on any non-403 response to cross-user flow execution attempts: `grep 'POST /api/v1/responses' langflow.log | grep -v '" 403 '` piped to an email alert.

Evidence: This step restores service — no volatile evidence destruction occurs at re-enablement, but before reviewing historical flow execution logs ensure archived copies from Step 2 are preserved. During the 30-day monitoring window, retain daily snapshots of the Langflow `flow_run` table and API gateway logs to a write-protected evidence store, as these provide the longitudinal record needed to confirm no post-patch IDOR exploitation recurs and to support any regulatory disclosure review of unauthorized AI workflow executions that occurred during the exposure window.

Step 5: Post-Incident — Conduct a review of all API endpoints in Langflow and any adjacent AI workflow platforms for similar missing object-level authorization checks. Map findings against NIST IR-4 (Incident Handling) to update the incident response playbook with IDOR-specific detection and containment steps. Evaluate whether AI workflow platforms in your environment are subject to periodic authorization audits. Reference CIS 6.1 and CIS 6.2 to verify that access granting and revoking processes account for AI agent and workflow platform permissions, not just traditional IT systems.

NIST Phase: Post Incident

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

Controls: NIST IR-4 — Incident Handling, NIST IR-8 — Incident Response Plan, NIST IR-5 — Incident Monitoring, CIS 6.1 (IG1/IG2/IG3) — Establish an Access Granting Process, CIS 6.2 (IG1/IG2/IG3) — Establish an Access Revoking Process, CIS 7.1 (IG1/IG2/IG3) — Establish and Maintain a Vulnerability Management Process, CIS 7.2 (IG1/IG2/IG3) — Establish and Maintain a Remediation Process

Compensating: For teams without a dedicated appsec function, enumerate all Langflow API routes from the application source or OpenAPI spec (`GET /docs` or `/openapi.json`) and manually test each authenticated endpoint with a second user's token attempting to access the first user's object IDs — document any non-403 responses as IDOR candidates. Use OWASP's BOLA/IDOR testing checklist as the review template. Add a Sigma rule to the detection library targeting POST /api/v1/responses with mismatched user-to-flow-owner pairs for any future Langflow deployment: `sigma convert -t splunk langflow_idor.yml`.

Evidence: No live-state evidence is at risk in this phase; however, compile the complete incident artifact package before closing the case: the pre-patch database export (flow, user, api_key, flow_run tables), archived API gateway logs covering the full exposure window, the pip freeze pre/post-patch manifests, and the authorization validation test results from Step 4. This package constitutes the evidentiary record for regulatory disclosure assessment, lessons-learned documentation, and playbook updates specific to IDOR vulnerabilities in AI workflow platforms.

Detection Guidance

Focus detection on the /api/v1/responses API endpoint. Parse application and API gateway logs for POST or GET requests to /api/v1/responses where the authenticated user identity (derived from the session token or API key) does not match the recorded owner of the flow_id parameter submitted in the request body or query string. A single account referencing multiple distinct flow IDs, particularly those owned by different users, is a strong indicator of exploitation. Behavioral indicators include: unusual spike in flow execution events from a single account, flow executions triggering on flows the user did not create, and API calls originating from IPs not associated with the account's historical access pattern. Per NIST AU-6, audit records should capture the authenticated user identity, source IP, flow ID requested, and execution outcome. CIS Controls v8 8.2 requires audit log collection to be enabled across the environment, confirm Langflow application logs are being shipped

to your SIEM. No public IOC hashes or network indicators are available from the provided source material at this time.

Framework Mappings

MITRE-ATTACK

- **T1565.001** — Stored Data Manipulation
- **T1078** — Valid Accounts

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **IR-5** — Incident Monitoring

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

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1565.001	Stored Data Manipulation	Impact
T1078	Valid Accounts	Defense-Evasion

Sources

Source	URL	Tier
vulncheck_kev	https://nvd.nist.gov/vuln/detail/CVE-2026-55255	T1
CVE-2026-55255 in langflow | Intelligent Application Security ...	https://www.kodemsecurity.com/cve-archive/cve-2026-55255	T3
CVE-2026-55255 - Vulnerability Details - OpenCVE	https://app.opencve.io/cve/CVE-2026-55255	T3

Source	URL	Tier
Endor Patches | CVE-2026-55255, Langflow: IDOR Vulnerability in ...	https://www.endorlabs.com/vulnerability/cve-2026-55255	**T3**
CVE-2026-55255 | Tenable®	https://www.tenable.com/cve/CVE-2026-55255	**T3**
CISA KEV	https://www.cisa.gov/known-exploited-vulnerabilities-catalog	**T1**