EXPLOITED VULNERABILITY ADVISORY: CVE-2026-64849

Summary

CVE-2026-64849 in MLflow MLflow is being actively exploited in the wild. MLflow is an open source AI engineering platform for agents, large language models, and machine learning models. Prior to 3.15.0, the unauthenticated POST /api/2.0/mlflow/webhooks/{id}/test endpoint calls _validate_webhook_url() in mlflow/utils/validation.py only for the original URL while mlflow/webhooks/delivery.py follows redirects and re-resolves the hostname without pinning the validated address, allowing attackers to reach internal or cloud metadata services and receive response_status and response_body. This issue is fixed in version 3.15.0. CISA requires federal remediation by 2026-09-02; under the CRA, exploitation of an affected product must be reported within 24 hours.

What we know

- **CVE:** CVE-2026-64849
- **Affected:** MLflow MLflow
- **Exploitation:** Actively exploited, added to CISA KEV on 2026-08-19
- **Severity:** CVSS 9.3 Critical, Network-exploitable, no privileges required, no user interaction (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:C/C:H/I:L/A:N)
- **Exploitation likelihood (EPSS):** 1%
- **Weakness:** CWE-918

Recommended action

Patch CVE-2026-64849 by 2026-09-02. If you cannot patch immediately, restrict network access to the affected service now (segment it, put it behind a VPN, or take it offline) until you can patch. If you ship an affected product, be ready to file the CRA 24-hour report the moment you find it exploited.

Sources: CISA Known Exploited Vulnerabilities, NIST NVD. Public data; an advisory of active exploitation, not an incident report. Verify applicability to your estate before acting.