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

3 September 2026

Advisory 239: Kestra Unauthenticated Remote Code Execution via Authentication Bypass (CVE-2026-49869).

Release Date: 2nd September 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, System/Network administrators, and Data/DevOps engineering teams that deploy or operate Kestra as a self-hosted workflow orchestration platform. This alert is intended to be understood by technical users and systems administrators.

What is it?

CVE-2026-49869 is a critical authentication-bypass vulnerability leading to unauthenticated remote code execution in Kestra, an open-source, event-driven orchestration platform for data, AI, and infrastructure workflows (27,000+ GitHub stars). Kestra's AuthenticationFilter uses a suffix match - checking whether a request path ends with "/configs" - to whitelist its public configuration endpoint from Basic Authentication, rather than validating the full path. Because the check only looks at the path's ending, an unauthenticated attacker can reach any API route that happens to end in "/configs", including routes for creating, executing, and deleting workflows, and for writing key-value secrets, none of which were intended to be exposed without credentials. Since Kestra ships with script-execution plugins (shell, Python, Node.js, and others) enabled by default, an attacker who reaches these endpoints can create and run an arbitrary workflow containing a malicious script task, achieving remote code execution as root inside the Kestra worker container without ever presenting a valid credential.

What are the systems affected?

The following version(s) are affected:

- Kestra before 1.0.45 – (Affected)
- Kestra 1.1.0 and later, before 1.3.21 – (Affected)

Kestra 1.0.45, and Kestra 1.3.21 and later – (Not affected, patched)

Kestra is typically self-hosted (via Docker, Helm chart, or a standalone JAR) rather than centrally auto-updated by the vendor, so this fix will not reach an existing deployment on its own.

What does this mean?

Typical attack flow:

1. **Reach the authentication-bypass endpoint**
    - An attacker sends an unauthenticated HTTP request to a reachable Kestra instance, targeting an API path ending in "/configs" - for example PUT /api/v1/main/flows/{namespace}/configs - which Kestra's Authentication Filter incorrectly whitelists based on the path's suffix alone, bypassing Basic Authentication entirely.
2. **Create and execute a malicious workflow**
    - Using the same authentication-bypass pattern, the attacker creates a workflow ("flow") containing a task that runs a shell, Python, or Node.js script - plugins Kestra enables by default - then triggers its execution via a similarly-bypassed endpoint (e.g. POST /api/v1/main/executions/{namespace}/configs), and retrieves the command output from the execution logs, achieving remote code execution as root inside the worker container without ever authenticating.

Attack vectors:

- A network-based, unauthenticated request against any internet- or network-reachable Kestra instance (typically listening on port 8080), requiring no prior access or account.
- No user interaction, no privileges, and no special access conditions are required (CVSS AV:N/AC:L/PR:N/UI:N) - a small number of crafted HTTP requests to documented, publicly-known endpoints is sufficient to achieve full compromise.

Successful exploitation may allow attackers to:

- Achieve remote code execution as root inside the Kestra worker container, and from there pivot into any systems, credentials, or data the worker has network access to.
- Perform server-side request forgery against internal network services and cloud metadata endpoints (AWS/GCP/Azure), potentially harvesting cloud credentials; overwrite or delete existing workflows and key-value secrets; and delete audit logs to cover evidence of the intrusion.

Mitigation process

CERTVU recommends the following:

1. Apply the vendor patch without delay

- Upgrade Kestra to version 1.0.45 or 1.3.21 (whichever branch applies to your deployment), where this vulnerability is fixed (GitHub Security Advisory GHSA-5vc5-wxxq-3fjx). Because Kestra is self-hosted, this requires a deliberate redeploy - confirm the running instance has actually reached a fixed version rather than assuming an update has applied.

2. Until patched, block network access to the instance from untrusted networks via firewall or reverse-proxy rules, since the vendor advisory documents no other workaround
3. Audit the estate for Kestra deployments
4. Treat any internet-reachable, unpatched instance as potentially compromised
5. Rotate credentials and secrets

Report suspected compromise to CERTVU at incident@cert.gov.vu or on telephone (678) 33380.

Reference

1. https://www.cve.org/CVERecord?id=CVE-2026-49869
2. https://github.com/kestra-io/kestra/security/advisories/GHSA-5vc5-wxxq-3fjx