

Artifactory Under Attack: In-the-Wild Exploitation of CVE-2026-42016, CVE-2026-42018 & CVE-2026-82329

✦ wiz.io/blog/artifactory-under-attack-in-the-wild-exploitation-of-cve-2026-42016-cve-2026-42018-cve-2026-82329

Shahar Dorfman, Sean Johnstone, Zohar Kaplan, Kurt Giacchino

September 10, 2026



Change log:

September 11, 2026 3PM UTC: Corrected the HTTP method for CVE-2026-82329 and the impacted fixed versions for CVE-2026-42018.

Wiz Research has identified active, in-the-wild exploitation of three critical and high-severity vulnerabilities affecting JFrog Artifactory: CVE-2026-42016, CVE-2026-42018, and CVE-2026-82329. Attackers are chaining these vulnerabilities to bypass authentication, escalate privileges, and gain administrative control over vulnerable Artifactory instances. Observed post-exploitation activity includes the creation of persistent administrator accounts, the deployment of malicious Groovy plugins for code execution, and the installation of Rust-based backdoors to establish persistence. This blogpost provides an analysis of the exploitation patterns observed, the impact on affected organizations, and actionable guidance for security teams to detect and remediate these threats. We will continue to update this blogpost as new information becomes available.

CVE-2026-42018: Exposure of an internal anonymous-user token

[CVE-2026-42018](#) is an improper-authentication vulnerability that may cause Artifactory to return an internal anonymous-user token to an unauthenticated requester, even when anonymous access is disabled. An attacker could use the exposed token to access resources available to the internal anonymous identity, potentially exposing sensitive artifacts or repository data.

CVE-2026-42016: Token scope validation flaw

[CVE-2026-42016](#) is a privilege-escalation vulnerability caused by insufficient token validation. Artifactory validates the token's signature and issuer but does not properly enforce its intended scope. As a result, an attacker with low-privileged access may be able to use a valid token to perform unauthorized actions and gain elevated privileges.

CVE-2026-82329: Unauthenticated access to administrative privileges

[CVE-2026-82329](#) is a critical authentication-bypass vulnerability affecting Artifactory under its default configuration. An unauthenticated attacker with network access to a vulnerable instance may be able to obtain administrative privileges, potentially gaining complete control over the Artifactory deployment and the artifacts, credentials, and integrations it manages.

What is the risk to cloud environments?

Our data indicates that 67% of organizations running JFrog Artifactory had at least one vulnerable instance when CVE-2026-42016 was first published on July 27. Similar levels were observed for CVE-2026-42018 (69% at publication on Aug 12) and CVE-2026-82329 (67% at publication on Aug 28).

Patching velocity has been slow for the lower-severity CVEs. As of six weeks after the first disclosure, 59% of organizations remain vulnerable to CVE-2026-42016, and CVE-2026-42018 has only declined from 69% to 62% over four weeks. However, CVE-2026-82329 has seen significantly faster remediation, dropping from 67% to 49% within two weeks of publication, likely due to its critical severity rating driving more urgent attention from security teams.

What evidence of exploitation has Wiz Research identified?

Wiz Research has confirmed in-the-wild exploitation of all three vulnerabilities across multiple environments.

CVE-2026-42018 and CVE-2026-42016 Exploitation

Between August 15 and September 8, 2026, we observed multiple actors chain CVE-2026-42018 and CVE-2026-42016 against self-hosted Artifactory instances. Across multiple cases we observed a custom Rust backdoor with C2 capabilities being dropped. Wiz Research is not aware of any prior public reporting of in-the-wild exploitation involving those two CVEs. Neither vulnerability grants administrative control on its own. CVE-2026-42018 exposes a token for the internal anonymous user, and CVE-2026-42016 lets that low-privileged token be escalated to admin scope. Together, the two can turn an unauthenticated request into an admin-scoped token in two steps.

Every exploitation followed a similar shape. An unauthenticated POST `/access/api/v1/aws/token/` with a trailing slash returned HTTP 200 with a JWT for the internal anonymous user, exploiting CVE-2026-42018. The actor then exchanged that JWT for an admin-scoped token through POST `/access/api/v1/tokens`, which returned HTTP 200 and exploited the CVE-2026-42016 scope-validation flaw. The escalated token kept the anonymous username but carried admin authority, so later requests appear with an actor of `token:anonymous`. In some instances, actors moved from the first request to a created admin account in under five minutes.

A request sequence looked like this:

- POST `/access/api/v1/aws/token/` 200 anonymous JWT for internal anonymous user (CVE-2026-42018)
- POST `/access/api/v1/tokens` 200 anonymous admin-scoped token (CVE-2026-42016)
- PUT `/api/security/users/<username>` or `/access/api/ui/users/<username>` 201 `token:anonymous` persistent admin account created

Since admin access is granted to the actor, post-exploitation varies. No single actor ran every step below. Across compromised Artifactory instances, we observed:

- **Persistent admin accounts** - PUT `/api/security/users/<username>` or `/access/api/ui/users/<username>`
- **Groovy plugin deployment** - malicious plugins installed through Artifactory's native plugin framework to gain arbitrary code execution on the server
- **Ad-hoc command execution** - shell commands run through the plugin endpoint **GET or POST** `/api/plugins/execute/<plugin>`, which include recon and file enumeration commands

- **Second-stage payload delivery** - a dropper fetched a binary over HTTP, wrote it to a world-writable path such as `/dev/shm`, `/tmp`, or `/var/tmp`, and established C2 communication
- **Webshell upload** - actors occasionally uploaded a script into a repository path to maintain follow-on access

CVE-2026-82329 Exploitation

This vulnerability has been reported by several vendors as being actively exploited in the wild, and has also been added to [CISA KEV](#). Between September 1 and September 8, 2026, we observed several threat actors carry out successful exploitations of CVE-2026-82329. Every initial exploitation followed the same pattern: an unauthenticated POST `/access/api/v1/registry/join` returning HTTP 200 or 201 with an admin-scoped token in the response body, followed by post-exploitation activities. Rather than a unified attack chain by a single threat actor, the following behaviors represent distinct patterns observed across various affected Artifactory environments:

- **Configuration exfiltration** - GET `/api/system/configuration`.
- **Persistent admin accounts** - PUT `/api/security/users/<username>` with `customData.artifactory_admin: true`.
- **Long-lived credentials** - token minting via `/access/api/v1/tokens`.
- **Cluster key theft** - several operators pulled the join key directly from `/access/api/v1/system/security/join_key`.
- **Enumeration** - Users, repos and tokens enumeration with GET `/access/api/repositories/<name>` or `/api/repositories/<name>`, GET `/access/api/security/users` or `/api/security/users` and GET `/access/api/security/tokens` or `/api/security/tokens`
- **Bring-your-own-Key** - In some cases we also observed attackers attaching their own SSH keys to the created users.

Most attacker-created accounts carry proof-of-concept boilerplate: usernames such as `Nxploited_[a-zA-z0-9]{3}`, `labadmin_<hex>`, `0xTerror`, and `svc_[a-zA-z0-9]{8}`. Some actors created more legitimate-looking accounts like: `jfrog-distribution`, `jfrog-insight`, `repo-service` and more.

How can Artifactory users detect exploitation?

CVE-2026-42018

The exploit sends `/access/api/v1/aws/token/`. The highest-confidence signature is behavioural: a 401 on the bare path followed by a 200 on any variant of it, from the same client, inside a short window. That pattern is an operator confirming the vulnerable variant before relying on it, and it is not something a normal client produces.

CVE-2026-42016

The detectable anomaly is entirely behavioural - a mismatch between who the identity is and what it is doing:

- The internal anonymous identity, or any low-privilege identity, minting tokens at `POST /access/api/v1/tokens` or `POST /artifactory/api/security/token`.
- A low-privilege identity enumerating users at `GET /access/api/v1/users` or `GET /artifactory/api/security/users/<username>`.
- A low-privilege identity reading or writing `/artifactory/api/plugins`.

CVE-2026-82329

Exploitation will trigger a POST request to `/access/api/v1/registry/join` with successful (200 or 201) status code. But this indication on its own is not enough to determine exploitation. Correlate those requests with known attacker actions, such as:

- `PUT /api/security/users/<username>` 201 persistent admin account created
- `GET /api/system/configuration` 200 Configuration extraction
- `GET /access/api/security/tokens` 200 tokens enumeration
- `GET /access/api/security/users` 200 users enumeration

Which products are affected?

CVE ID	Vulnerability Type	Impacted Versions
CVE-2026-82329	Authentication Bypass	7.161.0 > 7.161.19 7.146.0 > 7.146.36 7.133.0 > 7.133.28 7.125.0 > 7.125.19 7.117.0 > 7.117.27 7.111.4 > 7.111.21
CVE-2026-42018	Authentication Bypass	< 7.111.20; 7.117.0 -> 7.117.27; 7.125.0 -> 7.125.19; 7.133.0 -> 7.133.28; 7.146.0 -> 7.146.8
CVE-2026-42016	Privilege Escalation	Prior to 7.133.11

What actions should security teams take?

Organizations running JFrog Artifactory should identify affected instances and upgrade to a fixed version as soon as possible. Depending on the deployed release branch, fixed versions include 7.111.21, 7.117.28, 7.125.20, 7.133.29, 7.146.38, and 7.161.20 or later. In addition, customers should review logs and refer to the "How can Artifactory users detect exploitation?" section for guidance on identifying potential indicators of compromise.

Given that exploitation may be possible remotely without authentication under the default configuration, organizations should prioritize internet-accessible Artifactory instances and restrict network access to trusted users and systems where possible. Organizations should also review Artifactory authentication and administrative activity for unexpected privileged access.

How can Wiz help?

Wiz customers should refer to the [pre-built advisory](#) in the Wiz Threat Intel Center for actionable steps to investigate, remediate, and harden their environments. Wiz Defend customers benefit from [agentless scanning](#) of Artifactory logs. Wiz Research will continue to update our advisory and this blogpost as the situation develops.

If you suspect a JFrog Artifactory breach or any other cloud security incident, [contact Wiz CIRT](#) for incident response support.

Indicators of Compromise (IOCs)

Indicator	Description	First Seen	Last Seen
hxxp://log.gitclone[.]org:45678/smtp	Payload download after CVE-2026-42018/CVE-2026-42016 exploitation	2026-09-06	2026-09-08
hxxp://3.88.162[.]79:36789/smtp	Second load of payload after CVE-2026-42018/CVE-2026-42016 exploitation	2026-09-07	2026-09-08
513a907b69edffc3cb77a494da395178d21ef9bd	SHA1 of /tmp/.z payload	2026-09-06	2026-09-08
64.207.232[.]6:8443	C2 address	2026-09-08	2026-09-08
149.102.229[.]150	Actor IP exploiting CVE-2026-42018/CVE-2026-42016	2026-08-30	2026-08-30
186.247.79[.]240	Actor IP exploiting CVE-2026-42018	2026-09-02	2026-09-02
182.62.201[.]69	Actor IP exploiting CVE-2026-42018	2026-09-04	2026-09-04
146.19.216[.]120	Actor IP exploiting CVE-2026-82329	2026-09-01	2026-09-01
185.190.58[.]172	Actor IP exploiting CVE-2026-82329	2026-09-03	2026-09-03
45.61.176[.]88	Actor IP exploiting CVE-2026-82329	2026-09-02	2026-09-04

Indicator	Description	First Seen	Last Seen
223.144.227[.]110	Actor IP exploiting CVE-2026-82329	2026-09-04	2026-09-04
129.121.56[.]234	Actor IP exploiting CVE-2026-82329	2026-09-04	2026-09-04
16.54.250[.]190	Actor IP exploiting CVE-2026-82329	2026-09-04	2026-09-04
105.188.75[.]16	Actor IP exploiting CVE-2026-82329	2026-09-03	2026-09-03
103.124.165[.]42	Actor IP exploiting CVE-2026-82329	2026-09-03	2026-09-03
176.88.121[.]152	Actor IP exploiting CVE-2026-82329	2026-09-03	2026-09-03
155.254.120[.]23	Actor IP exploiting CVE-2026-82329	2026-09-02	2026-09-02
220.246.124[.]92	Actor IP exploiting CVE-2026-82329	2026-09-06	2026-09-07
15.157.64[.]113	Actor IP exploiting CVE-2026-82329	2026-09-04	2026-09-04
104.28.251[.]139	Actor IP exploiting CVE-2026-82329	2026-09-04	2026-09-04
137.184.111[.]69	Actor IP exploiting CVE-2026-82329	2026-09-02	2026-09-02
svc_[a-zA-z0-9]{8}	Malicious account regex created by threat actor		

Indicator	Description	First Seen	Last Seen
Nxploited_[a-zA-z0-9]{3}	Malicious account regex created by threat actor		
labadmin_[a-zA-z0-9]{10}	Malicious account regex created by threat actor		
jfrog-distribution	Malicious admin account created by threat actor		
backup-service	Malicious admin account created by threat actor		
repo-service	Malicious admin account created by threat actor		
jfrog-insight	Malicious admin account created by threat actor		
jfrog-mission-control	Malicious admin account created by threat actor		
jfrog-pipeline	Malicious admin account created by threat actor		
migration-tool	Malicious admin account created by threat actor		

Indicator	Description	First Seen	Last Seen
ldap_admin	Malicious admin account created by threat actor		
ldap_administrator	Malicious admin account created by threat actor		
0xterror	Malicious admin account created by threat actor		

- [CVE-2026-42018: Exposure of an internal anonymous-user token](#)
- [CVE-2026-42016: Token scope validation flaw](#)
- [CVE-2026-82329: Unauthenticated access to administrative privileges](#)
- [What is the risk to cloud environments?](#)
- [What evidence of exploitation has Wiz Research identified?](#)
- [How can Artifactory users detect exploitation?](#)
- [Which products are affected?](#)
- [What actions should security teams take?](#)
- [How can Wiz help?](#)
- [Indicators of Compromise \(IOCs\)](#)

Get a personalized demo

Ready to see Wiz in action?

"Best User Experience I have ever seen, provides full visibility to cloud workloads."

David EstlickCISO

"Wiz provides a single pane of glass to see what is going on in our cloud environments."

Blackstone

Adam FletcherChief Security Officer

"We know that if Wiz identifies something as critical, it actually is."

Greg PoniatowskiHead of Threat and Vulnerability Management

[Get a demo](#)