THE CYBER TOOLCHAIN

NO. 078

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 18, 2026 · EVERY DAY

The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 276 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 39 tools

▾ // THE BRIEF HIDE

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

Three unrelated things stand out. GitHub made workflow execution protections GA with default `pull_request_target` blocking, closing ctions path that leaks write-scoped tokens to fork PRs. Grype can

vim ? help

now surface matches it had quietly suppressed, so you audit the filter instead of trusting it. Aviatrix shipped FIPS 140-3 validated IPsec and L7 East-West inspection in one controller release.

PROTECT

Stop a fork pull request from running with your repo's write token

GitHub's workflow execution protections went GA with an API to set policy centrally and `pull_request_target` blocked by default — the trigger that runs untrusted fork code in a privileged context, and the root of most Actions supply-chain compromises. Arcjet's label validation is the same instinct one layer up: confirm a remote policy actually applied, rather than discovering weeks later it was failing open.

GitHub · Arcjet

PROTECT

See the vulnerabilities your scanner decided not to tell you about

Grype will now emit matches its own matchers dropped — distro-fixed records, known false positives — so you can check whether the suppression logic is hiding something that matters in your build rather than taking the clean report at face value. sslscan's new flag enumerates key exchange groups across every supported TLS version instead of only the negotiated one, which is how you find a weak group still offered on a fallback path.

Grype · sslscan

IDENTIFY

Keep the CMDB honest without a quarterly reconciliation project

Axonius 9.0.7 adds a CMDB Reconciliation workspace that produces an audit trail for the sync, turning the usual spreadsheet diff between discovered assets and the system of record into a repeatable job. runZero's Autonomous Discovery expands scan scope as it finds new networks, so the subnet nobody documented gets scanned without someone widening a target list by hand.

Axonius · runZero

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Hacktron

Feedback ⓘ

A rebuilt findings view, and fix PRs that close themselves NOTES ↗ RANK

Hacktron is an AI-powered security code-review platform. Our AI agents review your repositories for exploitable vulnerabilities and surface them as pull-request comments and dashboard findings.

> Hacktron rebuilt its findings table and detail view, extended fix-PR automation to Azure DevOps, added self-managed GitLab and GitHub Enterprise Server connections for appliances, and introduced appliance license controls including a `max_scans` cap and expiry warnings, alongside a string of findings-metadata and settings improvements.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 max_scans entitlement field on appliance licenses NEW

90

Adds a `max_scans` entitlement field to appliance licenses, capping the total number of full scans that may start over the life of the license; the cap is checked at scan start, including restarts. Set via `max_scans: 500` in the license configuration.

Cap the total scans an appliance license permits so a `max_scans` quota is enforced from the first scan attempt.

yaml

```yaml
max_scans: 500
```

— Named config field with example value and enforcement mechanism.

A rebuilt findings view, and fix PRs th…

02 exclude.authors bot PR trust rule for GitHub IMPROVED

88

Adds `exclude.authors` trust-rule support for bot pull requests on GitHub, so bot accounts like `renovate[bot]` are now reviewed alongside GitLab and Bitbucket, where this was already supported. Configured via the `exclude.authors` list in `.hacktron/config.yaml`.

Prevent a specific bot account from receiving AI fix PRs by adding it to the author exclusion list in your Hacktron config.

```yaml
exclude:
  authors:
    - renovate[bot]
    - my-internal-bot
```

— Named config key with runnable YAML example and cross-host scope.

A rebuilt findings view, and fix PRs th…

---

03 **Pin GitLab Whitebox Scan to tag or commit** NEW 75

Pins a GitLab Whitebox Scan to a specific tag or commit SHA instead of the tip of a branch, entered in the Tag or commit field when starting a new scan.

Target a specific release tag for a GitLab Whitebox Scan instead of scanning the current branch tip — useful for auditing a pinned release candidate.

> 📍 In the Hacktron web app, go to Whitebox Scans › New Scan, select your GitLab repository, then enter the tag name (e.g. `v2.4.1`) or commit SHA in the Tag or commit field before starting the scan.

— Exact UI field and example value given for running it.

A rebuilt findings view, and fix PRs th…

---

04 **Rebuilt findings table and detail drawer** IMPROVED 75

Rebuilds the findings table and detail drawer with a unified header tray offering Severity, Status, Remediate, Copy markdown, and Mark as duplicate actions, plus previous/next navigation and a full-page collapse control.

— Concrete UI actions and controls named, but no exact path.

A rebuilt findings view, and fix PRs th…

---

05 **License detail and expiry warnings for appliances** NEW 75

Adds license detail and expiry warnings to organization settings on appliances — showing licensee, expiry date, days remaining, seat count, and covered products — with a 30-day-out banner and scan blocking after the grace deadline.

— Concrete fields, banner timing, and enforcement mechanism described.

A rebuilt findings view, and fix PRs th…

---

06 **OWASP Top 10:2025 and CWE metadata on findings** NEW 65

Findings gain an OWASP Top 10:2025 category, settable from the overflow menu and visible/filterable in the findings table and drawer, plus a CWE identifier and plain-

language fix recommendation shown in the detail view where the scan engine supplies them.

— Named UI locations but no API or config surface.

A rebuilt findings view, and fix PRs th…

07 **Download scan context documents from triage header** NEW

60

Adds a download button in the triage header to pull a scan's context documents — threat model, architecture notes, and uploaded files — without leaving the findings view.

— Names the control and document types, no API given.

A rebuilt findings view, and fix PRs th…

08 **Personal settings consolidation** IMPROVED 60

Consolidates Account and Settings into a single Personal settings page with sections for Profile, Security, Connected accounts, Preferences, Sessions, and Organizations.

— All sections listed, described only as UI reorganization.

A rebuilt findings view, and fix PRs th…

THINNER COVERAGE BELOW

09 **Self-service API key management** IMPROVED 55

Allows any organization member to create, rotate, and revoke their own API keys from personal settings without owner or admin involvement.

— States before/after permission change, no endpoint given.

A rebuilt findings view, and fix PRs th…

10 **Self-managed GitLab connection to appliances** NEW 55

Connects self-managed GitLab instances to appliances for merge-request scanning, enabling the Reviews link, PR Review dashboard tabs, and source control settings.

— Names enabled surfaces but not the connection steps.

A rebuilt findings view, and fix PRs th…

11 **GitHub Enterprise Server connection via GitHub App** NEW

55

Connects GitHub Enterprise Server to appliances via a GitHub App installation instead of a personal access token.

— Auth mechanism named, no setup steps given. A rebuilt findings view, and fix PRs th…

12 **Azure DevOps remediation support** NEW 50

Supports Azure DevOps repositories for remediation automations — fix pull requests can now be opened and tracked on Azure DevOps alongside GitHub and GitLab.

— Names the platform but gives no setup steps or config.

A rebuilt findings view, and fix PRs th…

13 **Application filter scoped across PR Review dashboard** IMPROVED 50

Scopes the Application filter on the PR Review dashboard to every widget, including the security score card and the low-score banner.

— Names affected widgets but no steps to apply it.

A rebuilt findings view, and fix PRs th…

14 **Dedicated appliance onboarding flow** NEW 45

Provides a dedicated appliance onboarding flow — Welcome, Connect your code, You're all set — pre-filled for the configured Git host.

— Steps named but no config or command detail. A rebuilt findings view, and fix PRs th…

15 **Expanded pentest PDF exports** IMPROVED 45

Expands pentest PDF exports with more per-finding detail, display-order finding numbers, and syntax-highlighted code blocks.

— Lists improvements but no export command or endpoint.

A rebuilt findings view, and fix PRs th…

16 **Engine failure reason on scan detail page** NEW 30

Shows the engine-reported failure reason on the scan detail page when a scan fails.

— Bare one-line description with no further mechanism.

A rebuilt findings view, and fix PRs th…

17 **Severity badges in global search** NEW 30

Displays severity badges on findings in global search results.

— One-line UI addition, no further detail given. A rebuilt findings view, and fix PRs th…

A rebuilt findings view, and fix PRs that close themselves NOTES ↗ LATEST STABLE

Hacktron adds Azure DevOps remediation, Linear triage sync, appliance license controls, and a rebuilt findings view

└──▷ **USE IT**

Prevent a specific bot account from receiving AI fix PRs by adding it to the author exclusion list in your Hacktron config.

.hacktron/config.yaml

```
exclude:
  authors:
    - renovate[bot]
    - my-internal-bot
```

Cap the total scans an appliance license permits so a `max_scans` quota is enforced from the first scan attempt.

yaml

```
max_scans: 500
```

Target a specific release tag for a GitLab Whitebox Scan instead of scanning the current branch tip — useful for auditing a pinned release candidate.

> 📍 In the Hacktron web app, go to Whitebox Scans › New Scan, select your GitLab repository, then enter the tag name (e.g. `v2.4.1`) or commit SHA in the `Tag or commit` field before starting the scan.

Adds `exclude.authors` trust-rule support for bot pull requests on GitHub, so Dependabot, Renovate, and similar bots are now reviewed alongside GitLab and Bitbucket.

Adds `max_scans` entitlement field to appliance licenses, capping the total number of full scans that may start over the life of the license, checked at scan start including restarts.

Supports Azure DevOps repositories for remediation automations — fix pull requests can now be opened and tracked on Azure DevOps alongside GitHub and GitLab.

Pins a GitLab Whitebox Scan to a specific tag or commit SHA instead of the tip of a branch.

Adds OWASP Top 10:2025 category to findings — settable from the overflow menu, visible in the findings table and drawer, and filterable.

+13 more

WAS THIS USEFUL? 👍 👎

◆ Wireless & RF

RFID Research Group proxmark3

Feedback ⓘ v4.23346

NOTES ↗ RANK

Iceman Fork - Proxmark3

> Proxmark3's Iceman Fork now simulates a full DESFire EV1 card over USB or RF and gains matching dump/emulator tooling, alongside new MyKey/COGES decoding, a VIGIK key-recovery script, and a string of key-check and card-format decoding improvements.

└──▷ **WHAT SHIPPED · 13 FEATURES** most completely described first ? what's the number?

01 DESFire dump, emulator load/save and view NEW 95

`hf mfdes eload`, `hf mfdes esave` and `hf mfdes eview` load, save and view a DESFire card dump in emulator memory; `hf mfdes esave --keep` preserves applications and files a reader deleted rather than omitting them. `hf mfdes dump` walks the whole PICC and writes a `hf-mfdes-<UID>-dump.json` image, auto-loading any `hf-mfdes-<UID>-keys.json` found, and `hf mfdes view` prints a saved dump file with no device attached.

Dump a full DESFire PICC to a JSON image, auto-loading any previously recovered keys, for offline analysis.

```
$ hf mfdes dump
```

View a saved DESFire dump file without needing the device present — useful for reviewing card structure at a desk.

```
$ hf mfdes view
```

Save a DESFire emulator image while retaining applications and files a reader deleted during the session.

```
$ hf mfdes esave --keep
```

— Names every dump/emulator command and file naming convention v4.23346

02 DESFire EV1 card simulation NEW 93

`hf mfdes sim` now answers a wide range of DESFire EV1 commands from a card image in emulator memory: ISO 7816 `SELECT` by DF name or file ID, `SetConfiguration` (disable FormatPICC, switch to random UID, set a user-defined ATS), `ReadSignature`, `GetKeyVersion`, `ChangeKeySettings` (enciphered, master key only), `ChangeKey` (plain and XORed key frame, reflected in `hf mfdes esave` output), `ChangeFileSettings` (plain or enciphered depending on access right), `CreateFile` / `DeleteFile` for all five EV1 file flavours, `CreateApplication` / `DeleteApplication`, `FormatPICC`, and `GetDFNames` / `GetISOFileIDs`. `hf mfdes etest` drives the same simulation over USB with no RF, letting testers inject RndB and a random UID.

— Names every DESFire command answered and the RF-less test path v4.23346

---

03 **Card-format decoding in hf mf view** IMPROVED 80

`hf mf view` now shows which VIGIK fields are RSA signed, decodes Urmet Captiv card number fields and data regions, and decodes Hexact sector 9 and 11 payload cross-checked against sectors 0 and 15 and the UID; `hf mf view --selftest` also checks the Hexact payload decoder in both directions.

— Names each card format and the cross-check mechanism v4.23346

---

04 **Unified hf mfu chk key-check command** IMPROVED 70

`hf mfu cchk` and `hf mfu aeschk` are merged into a single `hf mfu chk` command that auto-detects UL-C vs UL-AES.

— Names old and new commands and the detection behaviour v4.23346

---

05 **MyKey/COGES key decoding for SRIX4K** NEW 65

`hf 14b view` now decodes MyKey / COGES keys stored on SRIX4K tags, and `hf 14b view --selftest` runs the MyKey parser's self-tests.

— Names the decode target and selftest flag but no format detail v4.23346

---

06 **Broader key checking in hf mfdes chk** IMPROVED 65

`hf mfdes chk` now checks every key number an application declares, not only those named in a file's access rights, and can run with no arguments, falling back to the bundled dictionary the way `hf mfdes detect` does.

— Describes scope change and new no-argument fallback behaviour v4.23346

THINNER COVERAGE BELOW

07 **Faster USB transfer via word-aligned NG framing** IMPROVED 55

NG reply framing now aligns the payload to a word boundary, increasing USB transfer speed from 630 to 789 kB/s on RDV4.

— Gives exact before/after throughput numbers and mechanism v4.23346

08 **VIGIK public key recovery script** NEW 50

Adds `vigik_recover_pk.py`, a Python script to recover public keys from RSA signatures.

— Names the script but not its inputs or usage v4.23346

09 **Named DESFire card-damage status codes** IMPROVED 50

Adds named DESFire status codes for the five codes that indicate a card damaged or disabled itself, replacing the previously undifferentiated raw `-20` value.

— Explains the change but not the five specific code names v4.23346

10 **hf 14b valid command removed** BREAKING 45

`hf 14b valid` has been removed, replaced by the new MyKey parser in `hf 14b view`.

— States removal and replacement but no migration steps v4.23346

11 **Persistent pm3 history writes** IMPROVED 45

`pm3 history` now writes after every command, so history survives a kill rather than being saved only on a clean exit.

— Explains behaviour change but names no config surface v4.23346

12 **Emulator memory size doubled** IMPROVED 45

Emulator memory size is raised from 4096 to 8192 bytes.

— Gives exact before/after numbers but no usage detail v4.23346

13 **Emulator memory reporting in hw status** IMPROVED 35

`hw status` now reports EMULATOR memory.

— One-line addition with no further detail v4.23346

v4.23346 NOTES ↗ LATEST STABLE

Proxmark3 v4.23346 adds full DESFire EV1 simulation, new dump/view/test commands, and unified UL key-check.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v4.23346
https://github.com/RfidResearchGroup/proxmark3.git
# already have the repo? check out this version:
$ git checkout v4.23346
```

└──▷ **TRY IT**

Dump a full DESFire PICC to a JSON image, auto-loading any previously recovered keys, for offline analysis.

```
$ hf mfdes dump
```

View a saved DESFire dump file without needing the device present — useful for reviewing card structure at a desk.

```
$ hf mfdes view
```

Save a DESFire emulator image while retaining applications and files a reader deleted during the session.

```
$ hf mfdes esave --keep
```

Adds `hf mfdes etest` — drives the DESFire simulation over USB without RF, with injectable RndB and random UID.

Adds `hf mfdes sim` — DESFire EV1 simulation now answers authentication, secure messaging, file reads, writes, and transactions from the card image in emulator memory.

Adds `hf mfdes sim` — answers ISO 7816 SELECT by DF name or file ID with real status words, so a reader that opens with a SELECT reaches the simulated card.

Adds `hf mfdes sim` — SetConfiguration is answered, so a reader can disable FormatPICC, switch the card to a random UID, or give it a user-defined ATS.

Adds `hf mfdes sim` — ReadSignature is answered for a card image of a generation that has one, refused the way a genuine EV1 refuses it.

+28 more

WAS THIS USEFUL?  

◆ Reconnaissance & OSINT

Maigret

Feedback ⓘ v0.6.6 NOTES ↗ CODE ↗ RANK

Collect a dossier on a person by username from 3000+ sites

> Maigret v0.6.6 expands its site coverage with new detection engines and site migrations, and adds batch lookups via `--input-file` alongside validation improvements to `--self-check` and `--submit`.

└──▷ **WHAT SHIPPED · 6 FEATURES** most completely described first ? what's the number?

01 **Status-code based check creation in --submit fallback** NEW

75

Adds `status_code` check creation to the `--submit` fallback: when a claimed profile returns 2xx and the unclaimed probe returns non-2xx, that signal is now preserved.

— Mechanism and before/after explained, no direct command v0.6.6

02 **New detection engines for forums, wikis and federated platforms**

NEW 70

Adds a Vanilla engine, a MediaWiki engine, a MediaWikiJson engine, a DiscourseJson engine, a Lemmy engine, and a shared Mastodon engine using the `accounts/lookup` urlProbe, expanding structured detection across forum, wiki, and federated platforms. Moves 395 Discourse sites to the DiscourseJson engine and 164 wikis to MediaWikiJson engines for more reliable profile detection.

— All engines and migration counts named, but internal to tool v0.6.6

03 **Batch username lookups via --input-file** NEW 65

Adds `--input-file` flag to read identifiers from a file, one per line, instead of listing them all on the command line.

— Flag and input format named, directly usable v0.6.6

04 **Custom DNS resolver for self-check and submit** IMPROVED

65

Adds `--dns-resolver` support inside `--self-check` and `--submit` so custom resolvers are respected during site validation.

— Names the flags and commands affected, no example v0.6.6

THINNER COVERAGE BELOW

05 **XenForo login-wall detection for Hebrew and Persian boards**

IMPROVED

55

XenForo engine now counts Hebrew and Persian login-wall notices as account presence, reducing false `available` reports on boards in those languages.

— Behavior change explained, no user-facing surface v0.6.6

06 **Windows standalone binary on release page** NEW

45

Attaches the Windows standalone binary (`maigret_standalone.exe`) directly to each release page.

— Simple distribution change with clear download target v0.6.6

v0.6.6 NOTES ↗ CODE ↗ LATEST STABLE

Maigret v0.6.6 doubles its site database to 5,200+ and adds `--input-file` for batch username lookups

└──▷ GET THIS VERSION

```
$ git clone --branch v0.6.6 https://github.com/soxoj/maigret.git
# already have the repo? check out this version:
$ git checkout v0.6.6
```

Adds `--input-file` flag to read identifiers from a file, one per line, instead of listing them all on the command line.

Adds `--dns-resolver` support inside `--self-check` and `--submit` so custom resolvers are respected during site validation.

Adds `status_code` check creation to the `--submit` fallback: when a claimed profile returns 2xx and the unclaimed probe returns non-2xx, that signal is now preserved.

Adds a shared Mastodon engine using the `accounts/lookup` urlProbe for Mastodon instance coverage.

Adds a Vanilla engine, a MediaWiki engine, a MediaWikiJson engine, a DiscourseJson engine, and a Lemmy engine, expanding structured detection across forum and wiki platforms.

+3 more

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ

Copilot impact dashboard now shows feature engagement (2026-09-17) BLOG ↗

RANK

Join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> GitHub shipped GA workflow execution protections for Actions with a new policy API and default `pull_request_target` blocking, made Ubuntu 26.04 runners generally available ahead of an `ubuntu-latest` migration, and expanded the Copilot usage metrics API with feature-engagement and agentic CLI customization data.

└──▷ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Ubuntu 26.04 runner GA and ubuntu-latest migration

BREAKING

90

Adds `ubuntu-26.04` and `ubuntu-26.04-arm` as fully supported `runs-on` values for GitHub Actions on x64 and arm64. The `ubuntu-latest` label migrates from Ubuntu 24.04 to Ubuntu 26.04, rolling out gradually between October 19 and November 19, 2026; workflows depending on tools or packages changed between the two versions may break automatically during that window unless pinned to `ubuntu-24.04` or explicitly set to `ubuntu-26.04` / `ubuntu-26.04-arm`.

Explicitly target the new Ubuntu 26.04 ARM runner before the `ubuntu-latest` migration completes, so arm64 CI jobs run on the GA image immediately.

yaml

```yaml
jobs:
  build:
    runs-on: ubuntu-26.04-arm
    steps:
      - uses: actions/checkout@v4
      - run: make test
```

Pin a workflow to `ubuntu-24.04` to avoid breakage during the `ubuntu-latest` migration while you validate your toolchain against Ubuntu 26.04.

```yaml
jobs:
  build:
    runs-on: ubuntu-24.04
    steps:
      - uses: actions/checkout@v4
      - run: make test
```

— Exact runs-on values, migration window and runnable YAML pins

Ubuntu 26 generally available and lates…

02 Workflow execution protections in Actions reach GA

BREAKING

82

A new REST API lets you create, read, update, and delete actor and event rules — including workflow path conditions — at the enterprise, organization, and repository level for policy-as-code governance. Adds workflow file targeting to scope rules to specific files (e.g. `deploy.yml`) instead of a whole repo, an Insights view to audit how rules evaluate and enforce (including in evaluate/shadow mode), and evaluate mode itself is now GA for previewing impact before enforcement. A default rule disabling `pull_request_target` on public repositories with no applicable event policy runs in evaluate mode now and auto-enforces on November 2, 2026; repos that need the trigger must explicitly allow it via Settings > Actions > Policies before then.

If your public repo legitimately depends on `pull_request_target`, explicitly allow it in an event policy scoped to specific workflow files to opt out of the default blocking rule before November 2, 2026.

> 📍 In the GitHub console, go to your repository › Settings › Actions › Policies, select the default `pull_request_target` rule, choose Allow for the specific workflow files that require the trigger, and save.

— Names API scope, UI path and enforcement date but not exact endpoints

Workflow execution protections in GitHu…

03 Agentic CLI customization metrics in Copilot usage API

NEW

78

Adds `totals_by_skill`, `totals_by_custom_agent`, `totals_by_mcp`, `totals_by_slash_cmd`, and `totals_by_plugin` arrays to enterprise and organization Copilot usage metrics API reports, each listing up to five most-active items with an `interaction_count`. Also adds `distinct_skill_use_count`, `distinct_custom_agent_use_count`, `distinct_mcp_use_count`, `distinct_slash_cmd_use_count`, and `distinct_plugin_use_count` fields counting how many different CLI customization items were used; all fields appear in

per-user and aggregate 1-day reports, per-user 28-day reports, and `day_totals` entries in aggregate 28-day reports.

— Names every field but no endpoint path given Agentic CLI customizations now in the u…

04 **Copilot impact dashboard feature engagement and phase reporting**

NEW 73

Adds a `copilot_feature_engagement` object to enterprise and organization 28-day aggregate reports, exposing active-user totals and a `totals_by_feature` breakdown covering code completion, agent edit, passive Copilot code review, active Copilot code review, Copilot cloud agent, Copilot CLI, and Copilot app. Also adds a `users_in_phase_28d` field to AI adoption phase reporting, exposing the full rolling 28-day phase population alongside the existing `total_engaged_users` count, so the impact dashboard can show how many active users engaged with each feature on at least two days in the period.

— Names all report fields but no example or endpoint call

Copilot impact dashboard now shows feat…

└──▷ **BREAKING ON UPGRADE**

! `ubuntu-latest` will point to Ubuntu 26.04 (up from 24.04) after the October 19 – November 19, 2026 rollout; workflows that depend on tools, packages, or versions removed or changed between Ubuntu 24.04 and 26.04 may break automatically during this window.

! On November 2, 2026, GitHub will automatically enforce a default rule disabling `pull_request_target` on public repositories that had no applicable event policy before general availability — workflows relying on `pull_request_target` in those repos will be blocked unless explicitly allow-listed via an Actions event policy.

Copilot impact dashboard now shows feature engagement (2026-09-17) BLOG ↗

LATEST STABLE

Copilot impact dashboard adds per-feature engagement counts and a new 28-day phase population field to enterprise and org reports.

Adds `copilot_feature_engagement` object to enterprise and organization 28-day aggregate reports, exposing active-user totals and per-feature engagement counts.

Adds `totals_by_feature` breakdown within `copilot_feature_engagement`, covering code completion, agent edit, passive Copilot code review, active Copilot code review, Copilot cloud agent, Copilot CLI, and Copilot app.

Adds `users_in_phase_28d` field to AI adoption phase reporting, exposing the full rolling 28-day phase population as of each report day alongside the existing `total_engaged_users` (active-that-day) count.

Copilot impact dashboard now surfaces how many active users engaged with each Copilot feature on at least two days during the 28-day period, giving enterprise administrators a view of which experiences are widely adopted.

Agentic CLI customizations now in the usage metrics API (2026-09-17) BLOG ↗

STABLE

GitHub Copilot usage metrics API gains agentic CLI fields for skills, custom agents, MCP servers, slash commands, and plugins.

Adds `totals_by_skill`, `totals_by_custom_agent`, `totals_by_mcp`, `totals_by_slash_cmd`, and `totals_by_plugin` arrays to enterprise and organization Copilot usage metrics API reports, each listing up to five most-active items with an `interaction_count`.

Adds `distinct_skill_use_count`, `distinct_custom_agent_use_count`, `distinct_mcp_use_count`, `distinct_slash_cmd_use_count`, and `distinct_plugin_use_count` fields to the same reports, counting how many different CLI customization items were used in a period.

New agentic CLI metrics appear in per-user and aggregate 1-day reports, per-user 28-day reports, and `day_totals` entries in aggregate 28-day reports.

Ubuntu 26 generally available and latest migration (2026-09-17) BLOG ↗

STABLE

Ubuntu 26.04 runner is GA for GitHub Actions; `ubuntu-latest` migrates from 24.04 to 26.04 starting October 19.

└──▷ **USE IT**

Explicitly target the new Ubuntu 26.04 ARM runner before the `ubuntu-latest` migration completes, so arm64 CI jobs run on the GA image immediately.

.github/workflows/ci.yml

```
jobs:
  build:
    runs-on: ubuntu-26.04-arm
    steps:
      - uses: actions/checkout@v4
      - run: make test
```

Pin a workflow to `ubuntu-24.04` to avoid breakage during the `ubuntu-latest` migration while you validate your toolchain against Ubuntu 26.04.

.github/workflows/ci.yml

```
jobs:
  build:
    runs-on: ubuntu-24.04
    steps:
      - uses: actions/checkout@v4
      - run: make test
```

Adds `ubuntu-26.04` and `ubuntu-26.04-arm` as fully supported `runs-on` values for GitHub Actions workflows on x64 and arm64.

`ubuntu-latest` label migrates from Ubuntu 24.04 to Ubuntu 26.04, rolling out gradually between October 19 and November 19, 2026.

└──▷ **BREAKING ON UPGRADE**

! `ubuntu-latest` will point to Ubuntu 26.04 (up from 24.04) after the October 19 – November 19, 2026 rollout; workflows that depend on tools, packages, or versions removed or changed between Ubuntu 24.04 and 26.04 may break automatically during this window.

`Workflow execution protections in GitHub Actions generally available (2026-09-17)`

BLOG ↗ STABLE

GitHub Actions workflow execution protections are GA, adding workflow-file scoping, an audit Insights view, and a full REST API for policy-as-code management.

└──▷ **HOW TO FIND IT**

If your public repo legitimately depends on `pull_request_target`, explicitly allow it in an event policy scoped to specific workflow files to opt out of the default blocking rule before November 2, 2026.

📍 In the GitHub console, go to your repository › Settings › Actions › Policies, select the default `pull_request_target` rule, choose `Allow` for the specific workflow files that require the trigger, and save.

New REST API for Actions execution protections lets you create, read, update, and delete actor and event rules — including workflow path conditions — at the enterprise, organization, and repository level, enabling policy-as-code governance at scale.

Adds workflow file targeting to scope execution protection rules to specific workflow files (e.g. `deploy.yml`) rather than an entire repository, so different policies can apply to different workflows in the same repo.

Adds an Insights view to audit how execution protection actions evaluate and enforce rules across your enterprise, organization, and repositories — including in evaluate (shadow) mode before enforcement.

Rolls out a default protection rule that disables `pull_request_target` events on public repositories that have no applicable event policy, running in evaluate mode initially and auto-enforcing on November 2, 2026.

Evaluate mode (shadow mode) is now generally available, letting you run execution protection rules without blocking workflows and preview which runs would be affected before enforcement.

└──▷ **BREAKING ON UPGRADE**

! On November 2, 2026, GitHub will automatically enforce a default rule disabling `pull_request_target` on public repositories that had no applicable event policy before general availability — workflows relying on `pull_request_target` in those repos will be blocked unless explicitly allow-listed via an Actions event policy.

WAS THIS USEFUL? 👍 👎

GitLab

Feedback ⓘ

GitLab Duo CLI takes a task from goal to done (2026-09-17) BLOG ↗ RANK

The intelligent orchestration platform for DevSecOps, enabling teams and agents to ship trusted software at enterprise scale.

> GitLab shipped a major expansion of MCP server tools spanning CI/CD, merge requests, repositories, work items and vulnerabilities, alongside a new `/goal` command for autonomous task completion in Duo CLI, three hosted open weight models for Duo Agent Platform, and a full set of AI credit governance, export and self-service consumption tools.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 Merge request MCP tool suite NEW

93

Adds `save_merge_request` to open and update merge requests, `list_merge_requests` supporting group-scope queries, and `get_merge_request`, which returns diffs, conflicts, or approvals via the `include` param, with diffs defaulting to stats-level detail unless `detail: full_patch` is passed. Adds `save_merge_request_review` covering the full review flow — line-level or summary comments, thread replies, discussion resolution, full review submission, Duo review trigger, and approve/unapprove — and `accept_merge_request`, which merges immediately by default, waits for checks with `strategy: merge_when_checks_pass`, or joins a merge train with `add_to_merge_train_when_checks_pass`.

Merge an MR only after all CI checks pass, rather than immediately, to give an agent a safe merge strategy in automated pipelines.

```json
Call the `accept_merge_request` MCP tool with:
{
  "merge_request_iid": 42,
  "project_id": 123,
  "strategy": "merge_when_checks_pass"
}
```

— Every tool and parameter named, with a runnable merge-strategy example.

New MCP tools help platform teams scale…

02 Usage data export with per-event detail and cost-center mapping

NEW

93

Adds `Export usage data` from the `GitLab Credits dashboard` or `GitLab Flex month overview`, delivering a background-generated export by email with a secure download link, covering up to 31 days across three files: a daily summary, a per-event file, and a manifest. The per-event file surfaces `Event ID`, `timestamp`, `event type`, `product and flow type`, `session identifier`, `User ID`, `Namespace ID`, `Project ID`, `credits consumed` (quantity and unit), `LLM call count`, `token totals`, `models used`, `Flex capability`, and `subscription name`, enabling per-team showback and session-level spike investigation. A one-time namespace-to-cost-center mapping lets these monthly per-event exports automatically reflect finance cost structure without manual reconciliation.

Pull a per-event export to identify which user and session drove last week's credit spike, then map namespaces to finance cost centers for automatic chargeback.

> 📍 From the GitLab Credits dashboard (or GitLab Flex month overview), select `Export usage data`. Choose a date range up to 31 days. The export arrives by email with a secure download link containing three files: a daily summary, a per-event file, and a manifest. In the per-event file, filter rows by User ID and date to isolate the session behind the spike, or group by Namespace ID to produce a department-level showback line.

— Every exported field named, plus a runnable export walkthrough.

`See who spent your AI credits and set f…`

---

03 **Credit caps, overrides, and usage alerts for admins** NEW

89

Adds the `Credit caps` page (at `Settings > GitLab Credits > Credit caps` for group Owners, or `Admin > GitLab Credits > Credit caps` on Self-Managed) to set a flat per-user GitLab Credits cap for all users in a billing period, with per-user overrides that can go above or below the flat cap, and a master toggle controlling whether overrides apply. Adds `GitLab Flex` capability-level controls letting admins mark each capability `Restricted`, `Usage cap`, or `Unlimited` against the shared reservation, so one capability's overage cannot drain credits needed by others. Adds email notifications to admins from the `GitLab Credits dashboard` at 50%, 80%, and 100% of credit usage, and again when the subscription cap is reached.

Set a default credit cap for all users and carve out a higher limit for engineers running agentic workloads — without opening the full budget to everyone.

> 📍 Group Owners: go to Settings › GitLab Credits › Credit caps for your top-level group. 1. Under `Flat user cap`, enter the credit limit (e.g. 50) and enable enforcement. 2. Select `Add override`, search for the engineers doing agentic work,

set their individual cap (e.g. 500), and save. 3. To revert an engineer to the flat cap, switch their override's enabled toggle off and save.

— Exact settings paths, thresholds, and a step-by-step example.

See who spent your AI credits and set f…

04 **Pipeline and job control MCP tools** NEW 82

Adds `save_pipeline` MCP tool to run, retry, cancel, or rename a CI/CD pipeline, and `get_job` MCP tool, which returns a job by default and, with `include: ["log"]`, streams the job trace paginated via `byte_offset` and `byte_limit` (512 KB per page).

Pull a failed CI job log in 512 KB pages so an agent can diagnose the build failure without fetching the full log at once.

json

```
Call the `get_job` MCP tool with:
{
  "job_id": 987654,
  "include": ["log"],
  "byte_offset": 0,
  "byte_limit": 524288
}
```

— Named tools and params with a runnable example; save_pipeline lacks its own sample call.

New MCP tools help platform teams scale…

05 **/goal command for autonomous task completion in Duo CLI** NEW 80

The `/goal` slash command in GitLab Duo CLI lets developers describe an outcome and a done-condition in a single prompt; the agent iterates locally without turn-by-turn supervision until the goal is met or a configured iteration limit is reached. A `/goal` run can be stopped at any point and a new goal issued mid-session, giving developers a clear intervention point without abandoning the workflow.

Delegate an open-ended fix to the agent and walk away — it iterates until all tests pass and CI is green, rather than waiting for you to re-prompt each step.

```
$ /goal Fix the failing tests in spec/models/user_spec.rb.
Continue until all tests pass locally and CI is green.
```

Hand off a broader clean-up task with an explicit done-condition and a human-review gate, so the agent stops only when the stated criteria are satisfied.

```
$ /goal Fix the failing tests. Ensure CI is fully green and the
fix is reviewed by Duo Developer.
```

— Full mechanism and runnable examples given; no iteration-limit number specified.

GitLab Duo CLI takes a task from goal t…

06 **Hosted open weight models in Duo Agent Platform** NEW

78

Adds three selectable hosted open weight models to GitLab Duo Agent Platform: `Kimi K3` (Moonshot AI) at 1.82 calls per credit, optimized for large codebases and long-running agent sessions; `GLM 5.3` (Z.ai) at 5 calls per credit, tuned for long-horizon coding tasks with extended context accuracy; and `MiniMax M3` (MiniMax) at 8 calls per credit, designed for high-volume, multi-step agentic work at scale. All three are served through Fireworks AI under a zero data retention policy that discards model input and output immediately after each response. Top-level group owners can set a per-feature default model and curate which models teams can select from, with settings applying consistently across all child groups and projects.

— Exact per-credit rates and provider named; no UI path for selection given.

Optimize your team's price-performance …

07 **Created trigger for merge requests in Duo Agent Platform** NEW

73

Adds a `Created` trigger action for merge requests in GitLab Duo Agent Platform, configurable under `AI > Triggers` in project settings, so flows or external agents run the moment an MR is opened and diffed.

Trigger an automated first-pass review flow the moment a contributor opens a merge request, by enabling the new `Created` trigger in your project.

> 📍 In your project, go to AI › Triggers and add the `Created` trigger action for merge requests, then enable your flow to run on that event.

— Exact settings path and a step-by-step example to enable it.

New MCP tools help platform teams scale…

08 **Vulnerability management MCP tools** NEW

69

Adds `list_vulnerabilities` (Ultimate only) filterable by severity and report type with cursor pagination, `get_vulnerability` for full single-vulnerability detail, and `save_vulnerability` (Ultimate only), which consolidates dismiss, confirm, revert to detected, update severity, and create a linked issue into a single tool.

— Tier restriction and consolidated actions named, but no sample call.

New MCP tools help platform teams scale…

09 **Self-service credit consumption view for developers** NEW

68

Adds a per-developer `GitLab Credits` view under `User settings` showing a daily consumption chart with product filter and date range controls, scoped to the individual user's own events.

Let a developer check their own AI credit consumption without needing admin access.

> 📍 As any member of the top-level group, navigate to User settings › GitLab Credits. Use the product filter and date range controls to view your daily consumption chart scoped to your own events.

— Exact settings path and controls, with a walkthrough example.

See who spent your AI credits and set f…

10 **MCP tool governance for read/write actions** IMPROVED

63

Extends MCP server tool governance to GitLab MCP server tools alongside GitLab Duo Agent Platform tools in the same group and project settings, with a per-tool mode: read-only tools default to `Always allow`, write/delete tools default to `Always ask`.

— Default modes named but no exact settings navigation given.

New MCP tools help platform teams scale…

11 **Work item and note MCP tools** NEW

61

Adds `get_work_item`, `list_work_items`, and `save_work_item` MCP tools covering issues, epics, and all supported work item types, superseding `get_issue` and `create_issue`. Adds `save_note` to comment on a work item or merge request and reply inside an existing thread, superseding `create_merge_request_note` and `create_workitem_note`.

— Names superseded tools but gives no call example.

New MCP tools help platform teams scale…

THINNER COVERAGE BELOW

12 **Repository interaction MCP tools** NEW

52

Adds `list_repository_tree`, `list_branches`, `list_tags`, `list_releases`, `list_commits`, `get_commit`, `add_commit`, and `fork_repository` MCP tools for full repository interaction.

— All eight tool names given but no usage example or parameters.

New MCP tools help platform teams scale…

13 **Promotional credit draw-down order change** IMPROVED

50

Changes the included promotional credit draw-down order: per-user included promotional credits (12 per Premium seat, 24 per Ultimate seat per month) are now applied before any shared evaluation credits are consumed.

— Exact seat allowances named but no action a user needs to take.

See who spent your AI credits and set f…

14 **GitLab Duo Slack agent (announced)** NEW

43

GitLab announced a forthcoming Duo Slack agent that will let developers mention `@GitLab` in any channel or thread to ask questions, trigger agent flows, or create GitLab issues conversationally, with progress streamed back into the thread in real time.

— Announcement only — not yet shipped, no config or usage path given.

GitLab Duo CLI takes a task from goal t…

15 **Project and user lookup MCP tools** NEW

42

Adds `get_project`, `list_projects`, `list_project_members`, and `get_user` MCP tools for project and user lookups.

— Tool names listed but no parameters or example calls.

New MCP tools help platform teams scale…

GitLab Duo CLI takes a task from goal to done (2026-09-17) BLOG ↗ STABLE

GitLab Duo CLI adds `/goal` slash command for autonomous, goal-driven task completion without turn-by-turn prompting.

└──▷ **TRY IT**

Delegate an open-ended fix to the agent and walk away — it iterates until all tests pass and CI is green, rather than waiting for you to re-prompt each step.

```
$ /goal Fix the failing tests in spec/models/user_spec.rb. Continue until all tests pass locally and CI is green.
```

Hand off a broader clean-up task with an explicit done-condition and a human-review gate, so the agent stops only when the stated criteria are satisfied.

```
$ /goal Fix the failing tests. Ensure CI is fully green and the fix is reviewed by Duo Developer.
```

Adds the `/goal` slash command to GitLab Duo CLI, letting developers describe an outcome and a done-condition in a single prompt; the agent iterates locally until the goal is met or the configured iteration limit is reached, with no turn-by-turn supervision required.

Supports stopping a `/goal` run at any point and issuing a new goal mid-session, giving developers a clear intervention point without abandoning the entire workflow.

Announces a forthcoming GitLab Duo Slack agent that will let developers mention `@GitLab` in any channel or thread to ask questions, trigger agent flows, or create GitLab issues conversationally, with progress streamed back into the thread in real time.

New MCP tools help platform teams scale automation safely (2026-09-17) BLOG ↗

STABLE

GitLab 19.4 adds extensive MCP server tools for CI/CD, MRs, vulnerabilities, and unified agent governance

└──▷ **USE IT**

Pull a failed CI job log in 512 KB pages so an agent can diagnose the build failure without fetching the full log at once.

```json
Call the `get_job` MCP tool with:
{
  "job_id": 987654,
  "include": ["log"],
  "byte_offset": 0,
  "byte_limit": 524288
}
```

Merge an MR only after all CI checks pass, rather than immediately, to give an agent a safe merge strategy in automated pipelines.

```json
Call the `accept_merge_request` MCP tool with:
{
  "merge_request_iid": 42,
  "project_id": 123,
  "strategy": "merge_when_checks_pass"
}
```

Trigger an automated first-pass review flow the moment a contributor opens a merge request, by enabling the new `Created` trigger in your project.

📍 In your project, go to AI › Triggers and add the `Created` trigger action for merge requests, then enable your flow to run on that event.

Adds `save_pipeline` MCP tool to run, retry, cancel, or rename a CI/CD pipeline.

Adds `get_job` MCP tool that returns a job by default, and with `include: ["log"]` streams the job trace paginated via `byte_offset` and `byte_limit` (512 KB per page).

Adds `save_merge_request` MCP tool to open and update merge requests.

Adds `get_merge_request` MCP tool that returns diffs, conflicts, or approvals via the `include` param, with diffs defaulting to stats-level detail unless `detail: full_patch` is passed.

Adds `list_merge_requests` MCP tool supporting group-scope queries.

`+10 more`

`Optimize your team's price-performance with hosted open weight models (2026-09-17)`

BLOG ↗ STABLE

GitLab Duo Agent Platform adds three hosted open weight models — Kimi K3, GLM 5.3, and MiniMax M3 — offering up to 4x more calls per GitLab Credit.

Adds `Kimi K3` (Moonshot AI) as a selectable hosted open weight model in GitLab Duo Agent Platform, at 1.82 calls per credit, optimized for large codebases and long-running agent sessions.

Adds `GLM 5.3` (Z.ai) as a selectable hosted open weight model in GitLab Duo Agent Platform, at 5 calls per credit, tuned for long-horizon coding tasks with extended context accuracy.

Adds `MiniMax M3` (MiniMax) as a selectable hosted open weight model in GitLab Duo Agent Platform, at 8 calls per credit, designed for high-volume, multi-step agentic work at scale.

Top-level group owners can now set a per-feature default model and curate which models teams can select from, with settings applying consistently across all child groups and projects.

All three new open weight models are served through Fireworks AI under a zero data retention policy, with model input and output discarded immediately after each response.

`See who spent your AI credits and set fair caps per team (2026-09-17)` BLOG ↗

LATEST STABLE

GitLab 19.4 adds per-user AI credit caps, per-event usage exports, and a developer self-service consumption view

└──▷ **HOW TO FIND IT**

Set a default credit cap for all users and carve out a higher limit for engineers running agentic workloads — without opening the full budget to everyone.

> 📍 Group Owners: go to Settings › GitLab Credits › Credit caps for your top-level group. 1. Under `Flat user cap`, enter the credit limit (e.g. 50) and enable enforcement. 2. Select `Add override`, search for the engineers doing agentic work, set their individual cap (e.g. 500), and save. 3. To revert an engineer to the flat cap, switch their override's enabled toggle off and save.

Pull a per-event export to identify which user and session drove last week's credit spike, then map namespaces to finance cost centers for automatic chargeback.

> 📍 From the GitLab Credits dashboard (or GitLab Flex month overview), select `Export usage data`. Choose a date range up to 31 days. The export arrives by email with a secure download link containing three files: a daily summary, a per-event file, and a manifest. In the per-event file, filter rows by User ID and date to isolate the session behind the spike, or group by Namespace ID to produce a department-level showback line.

Let a developer check their own AI credit consumption without needing admin access.

> 📍 As any member of the top-level group, navigate to User settings › GitLab Credits. Use the product filter and date range controls to view your daily consumption chart scoped to your own events.

Adds the `Credit caps` page (at `Settings > GitLab Credits > Credit caps` for group Owners, or `Admin > GitLab Credits > Credit caps` on Self-Managed) to set a flat per-user GitLab Credits cap for all users in a billing period, with per-user overrides that can go above or below the flat cap, and a master toggle that controls whether overrides apply.

Adds `Export usage data` from the `GitLab Credits dashboard` or `GitLab Flex month overview`, delivering a background-generated export by email with a secure download link — up to 31 days, three files per export: a daily summary, a new per-event file, and a manifest.

The new per-event export file surfaces `Event ID`, `timestamp`, `event type`, `product and flow type`, `session identifier`, `User ID`, `Namespace ID`, `Project ID`, `credits consumed` (quantity and unit), `LLM call count`, `token totals`, `models used`, `Flex capability`, and `subscription name` — enabling per-team showback and session-level spike investigation.

> Adds a per-developer `GitLab Credits` view under `User settings` showing a daily consumption chart with product filter and date range controls, scoped to the individual user's own events.
>
> Adds email notifications to admins from the `GitLab Credits dashboard` at 50%, 80%, and 100% of credit usage, and again when the subscription cap is reached.
>
> `+3 more`

WAS THIS USEFUL?  

SonarSource SonarQube

Feedback ⓘ announced · Sep 15, 2026

BLOG ↗ RANK

SonarQube is the independent AI code verification platform trusted by 7M+ developers & organizations to help build secure, reliable software with confidence.

> SonarQube expanded automatic analysis to GitLab with zero configuration and added native R language support with secrets detection, coverage import, and rule-based linting.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 Zero-configuration automatic analysis for GitLab NEW

75

SonarQube Cloud extends its zero-configuration automatic analysis — previously exclusive to GitHub and Azure DevOps — to GitLab, scanning every merge request and main-branch commit without editing `.gitlab-ci.yml`. Coverage spans 20+ languages including C#, Java, and Python; onboarding is done by linking a GitLab account with repository and metadata read permissions and using Import repositories in SonarQube Cloud.

Onboard a GitLab repository to SonarQube Cloud for instant, pipeline-free scanning on every merge request.

📍 1. Sign in to SonarQube Cloud and link your GitLab account (ensure your access token has repository and metadata read permissions). 2. Go to Import repositories and select the GitLab repositories you want to analyze. 3. For eligible repositories, automatic analysis starts immediately — no further configuration needed. 4. Results appear on the default branch and active merge requests shortly after import.

— Names mechanism, language scope, and exact onboarding steps launch-20260915-5d8d3f1f

02 **Native R language support** NEW 70

SonarQube adds analysis of R source files and embedded R code inside R Markdown documents, with 82 rules covering key sections of the lintr rule set plus additional Sonar-developed rules. It imports lintr problem reports directly into SonarQube's centralized findings view, adds Cobertura test coverage import for R projects, and extends secrets detection, code metrics, syntax highlighting, and copy-paste duplication analysis to R. Existing quality profiles, quality gates, PR decoration, and branch analysis now apply to R codebases on both Cloud and Server.

— Names rule count, lintr and Cobertura import, but no config steps

launch-20260915-f7587a71

announced · Sep 15, 2026 BLOG ↗

SonarQube Cloud adds zero-configuration automatic analysis for GitLab, scanning every merge request and main-branch commit without .gitlab-ci.yml edits.

└──▷ **HOW TO FIND IT**

Onboard a GitLab repository to SonarQube Cloud for instant, pipeline-free scanning on every merge request.

📍 1. Sign in to SonarQube Cloud and link your GitLab account (ensure your access token has repository and metadata read permissions). 2. Go to Import repositories and select the GitLab repositories you want to analyze. 3. For eligible repositories, automatic analysis starts immediately — no further configuration needed. 4. Results appear on the default branch and active merge requests shortly after import.

Extends zero-configuration automatic analysis — previously exclusive to GitHub and Azure DevOps — to GitLab, covering 20+ languages including C#, Java, and Python.

announced · Sep 15, 2026 BLOG ↗

SonarQube adds native R language support with 82 rules, secrets detection, lintr import, and Cobertura coverage on Cloud and Server.

Analyzes R source files and embedded R code inside R Markdown documents with 82 rules covering key sections of the lintr rule set plus additional Sonar-developed rules.

Imports lintr problem reports directly into SonarQube, surfacing existing findings in the centralized view alongside other languages.

Adds Cobertura test coverage import for R projects.

Extends secrets detection, code metrics, syntax highlighting, and copy-paste duplication analysis to R.

> Applies existing quality profiles, quality gates, PR decoration, and branch analysis to R codebases.

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Sep 18, 2026 DOCS ↗ BLOG ↗ RANK

The official changelog for the Snyk AI Trust Platform. Discover what's new in our security intelligence, AI-powered tools, and latest product enhancements.

> Snyk's CLI and API & Web platform gained internal CA trust support for DAST scans and clearer error reporting for unreadable policy files.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 **Custom internal CA upload for DAST scans** NEW 62

Snyk API & Web now supports uploading your organization's internal Root CA certificate, so certificates it issues are trusted during DAST scans alongside the public trust store. To use it, go to organisation settings in the Snyk Web UI, navigate to Add a custom CA certificate, and upload your Root CA file so subsequent scans verify certificates against it.

Upload your organisation's internal Root CA so DAST scans trust internally-issued certificates and avoid false positives on internal or internet-facing targets.

> 📍 In the Snyk Web UI, go to your organisation settings and navigate to Add a custom CA certificate, then upload your Root CA file so every subsequent scan verifies certificates against it alongside the public trust store.

— Names UI path and purpose but no CLI/API surface 2026-09-18-9290b991

02 **Descriptive error for unreadable .snyk policy files** IMPROVED 60

`snyk test` now reports an unreadable `.snyk` policy file as error code `SNYK-POLICY-0002` with a descriptive message instead of an unspecified error.

— Exact command and error code named, no further mechanism 2026-09-18-9290b991

changelog · Sep 18, 2026 DOCS ↗ BLOG ↗ LATEST STABLE

Snyk CLI v1.1307.3 adds per-project .snyk policy resolution; API & Web gains custom internal CA upload for DAST scans.

└──▷ **HOW TO FIND IT**

Upload your organisation's internal Root CA so DAST scans trust internally-issued certificates and avoid false positives on internal or internet-facing targets.

> 📍 In the Snyk Web UI, go to your organisation settings and navigate to Add a custom CA certificate, then upload your Root CA file so every subsequent scan verifies

certificates against it alongside the public trust store.

Adds support for uploading your organization's internal Root CA certificate in Snyk API & Web, so certificates it issues are trusted during DAST scans alongside the public trust store.

`snyk test` now reports an unreadable `.snyk` policy file as error code `SNYK-POLICY-0002` with a descriptive message instead of an unspecified error.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

sslscan

Feedback ⓘ 2.2.3 NOTES ↗ RANK

sslscan tests SSL/TLS enabled services to discover supported cipher suites

> sslscan 2.2.3 adds a flag to enumerate key exchange groups across every supported TLS protocol version instead of only the negotiated one.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **Scan key exchange groups across all TLS protocols** NEW

80

The new `--all-groups` flag scans for key exchange groups on all supported TLS protocols, not just the one negotiated in the handshake, e.g. `sslscan --all-groups <target>:443`.

Enumerate supported key exchange groups across all TLS protocol versions on a target to catch group misconfigurations not visible in a single handshake.

```
$ sslscan --all-groups <target>:443
```

— Named flag with runnable example, but limited mechanism detail. 2.2.3

2.2.3 NOTES ↗ LATEST STABLE

sslscan 2.2.3 adds `--all-groups` to enumerate key exchange groups across all supported TLS protocols.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 2.2.3 https://github.com/rbsec/sslscan.git
# already have the repo? check out this version:
$ git checkout 2.2.3
```

└──▷ **TRY IT**

Enumerate supported key exchange groups across all TLS protocol versions on a target to catch group misconfigurations not visible in a single handshake.

```
$ sslscan --all-groups <target>:443
```

WAS THIS USEFUL?  

Invicti

Feedback ⓘ Release 20260917 (2026-09-17) BLOG ↗ RANK

Official documentation for Invicti AppSec Platform, Invicti Platform, and Invicti Enterprise — web application security testing, DAST, API security, and vulnerability management.

> Invicti added a Website Discovery REST API for pulling discovered websites into external workflows, surfaced the source location of auto-detected API specifications in API Discovery, and expanded filtering options on the Website Discovery list.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Extended filters on Website Discovery list** IMPROVED 55

The Website Discovery list can now be filtered by discovery date, with exact match and exclusion options on text filters, and by multiple risk score values simultaneously.

– Names specific filter types a user can apply directly Release 20260917 (2026-09-17)

02 **Website Discovery REST API endpoint** NEW 45

A new REST API endpoint retrieves discovered websites, allowing discovery results to be integrated into external workflows and other cyber tooling.

– Names the API but no exact path or method given Release 20260917 (2026-09-17)

03 **API specification source location in API Discovery** IMPROVED 45

API Discovery now shows the source location of each API specification found during a scan, so users can see which endpoints were auto-detected and from where.

– Describes the change but not the exact UI path Release 20260917 (2026-09-17)

└──▷ **ALSO FROM THESE RELEASES**

Release 20260917 (2026-09-17) BLOG ↗ LATEST STABLE

Invicti adds a Website Discovery REST API and exposes API specification source locations in API Discovery.

Adds a Website Discovery REST API endpoint to retrieve discovered websites and integrate discovery results into external workflows and cyber tooling.

Shows the source location of each API specification found during a scan in API

> Discovery, so you can see which endpoints were auto-detected and from where.
>
> Enables filtering the Website Discovery list by discovery date, exact match and exclusion on text filters, and multiple risk score values simultaneously.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Volatility 3

Feedback ⓘ v2.28.2 NOTES ↗ CODE ↗ RANK

Volatility 3.0 development

> Volatility 3's v2.28.2 release adds a Mermaid tree renderer for embeddable diagrams and reworks Linux malfind's dirty-page dumping flags, alongside smaller improvements to layer validation and output handling.

└──▷ **WHAT SHIPPED · 5 FEATURES** most completely described first ? what's the number?

01 Mermaid renderer for tree visualization NEW

85

Adds a new `mermaid` renderer, selected with `--renderer mermaid`, that visualizes tree-structured plugin output such as process trees and VAD trees as Mermaid diagrams embeddable directly in Markdown and GitHub.

Visualize the Windows process tree as a Mermaid diagram to paste directly into a GitHub incident-report or Markdown wiki page.

```
$ vol -f /cases/memory.raw --renderer mermaid windows.pstree
```

— Names flag and includes a runnable command example v2.28.2

02 Linux malfind dirty-page dumping flags BREAKING

85

Replaces `--dump-page` and `--dump-size` with `--show-all-dirty-pages` and `--hexdump-size` in Linux `malfind`, giving finer control over dirty-page enumeration output; scripts or workflows using the old flags will break.

— Exact old and new flag names given, no example v2.28.2

THINNER COVERAGE BELOW

03 Banner redirected to stderr for structured output IMPROVED

55

Redirects the Volatility banner to stderr when a structured-output renderer (such as `mermaid`, CSV, or JSONL) is in use, preventing banner text from corrupting piped output.

— Names affected renderers, no exact command shown v2.28.2

04 Arrow/Parquet renderer removed from standalone EXE BREAKING

55

Arrow/Parquet output renderer is removed from the standalone executable build; users relying on Arrow output in the EXE must switch to the pip-installed package.

— Names affected format and migration path v2.28.2

05 **Windows Intel layer validation via translation checks**

IMPROVED 35

Improves Windows Intel layer validation by verifying layers via translation checks, increasing accuracy of memory layer identification.

— Mechanism named but no concrete usage detail v2.28.2

└──▷ **BREAKING ON UPGRADE**

! Linux `malfind` option `--dump-page` is replaced by `--show-all-dirty-pages`; scripts or workflows using `--dump-page` will break.

! Linux `malfind` option `--dump-size` is replaced by `--hexdump-size`; scripts or workflows using `--dump-size` will break.

! Arrow/Parquet output renderer is removed from the standalone executable build; users relying on Arrow output in the EXE must switch to the pip-installed package.

v2.28.2 NOTES ↗ CODE ↗ LATEST STABLE

Volatility 3 v2.28.2 adds a Mermaid renderer for tree visualization and expands Linux malfind's VMA dumping options.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.28.2
https://github.com/volatilityfoundation/volatility3.git
# already have the repo? check out this version:
$ git checkout v2.28.2
```

└──▷ **TRY IT**

Visualize the Windows process tree as a Mermaid diagram to paste directly into a GitHub incident-report or Markdown wiki page.

```
$ vol -f /cases/memory.raw --renderer mermaid windows.pstree
```

Adds `mermaid` renderer (`--renderer mermaid`) for visualizing tree-structured plugin output (process trees, VAD trees) as Mermaid diagrams embeddable in Markdown and GitHub.

Replaces `--dump-page` and `--dump-size` with `--show-all-dirty-pages` and `--hexdump-size` in Linux `malfind`, giving finer control over dirty-page enumeration

output.

Improves Windows Intel layer validation by verifying layers via translation checks, increasing accuracy of memory layer identification.

Redirects the Volatility banner to stderr when a structured-output renderer (such as `mermaid`, CSV, or JSONL) is in use, preventing banner text from corrupting piped output.

└──▷ **BREAKING ON UPGRADE**

- ! Linux `malfind` option `--dump-page` is replaced by `--show-all-dirty-pages`; scripts or workflows using `--dump-page` will break.
- ! Linux `malfind` option `--dump-size` is replaced by `--hexdump-size`; scripts or workflows using `--dump-size` will break.
- ! Arrow/Parquet output renderer is removed from the standalone executable build; users relying on Arrow output in the EXE must switch to the pip-installed package.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Zeek

Feedback ⓘ announced · Sep 17, 2026 BLOG ↗ RANK

Zeek is a powerful network analysis framework that is much different from the typical IDS you may know.

> Zeek 9 switches cluster communications to ZeroMQ by default, adds a declarative `zeek.conf` for systemd-based cluster orchestration, native encapsulated-traffic ingestion for cloud deployments, and a wide set of new log fields and analyzer capabilities — alongside breaking SHA256 upgrades to `ssh.log` and `kerberos.log` fingerprint/ticket fields and a Spicy 1.17.0 performance upgrade.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 **Software framework caching tunables** NEW 90

Adds `Software::parse_cache_interval`, `Software::found_cache_interval`, and `Software::max_software_cache_size` tunables to the software framework for finer control over version-detection caching, letting operators cap cache size and intervals to reduce memory use and log churn in high-volume environments.

Cap the software framework's detection cache to limit memory use and reduce noisy software-version log churn in high-volume environments.

zeek

```
redef Software::max_software_cache_size = 5000;
redef Software::parse_cache_interval = 1min;
redef Software::found_cache_interval = 5min;
```

— Named tunables paired with a runnable config example. launch-20260917-0d171be8

02 **Encrypted Client Hello (ECH) support in TLS analyzer** NEW 85

Adds a new `ssl_extension_encrypted_client_hello` event in the TLS analyzer to report on Encrypted Client Hello (ECH), and adds `ech_kdf_id` and `ech_aead_id` fields to `ssl.log` — reporting HPKE KDF and AEAD identifiers from the TLS 1.3 ECH extension — via the `protocols/ssl/ssl-log-ext.zeek` policy script.

Enable ECH extension logging in ssl.log to capture HPKE KDF and AEAD identifiers from TLS 1.3 Encrypted Client Hello handshakes.

```
$ @load protocols/ssl/ssl-log-ext
```

— Named event, fields, and policy script with a load example. launch-20260917-0d171be8

03 **Kerberos ticket fields upgraded to SHA256** BREAKING

85

Replaces the MD5-hashed `auth_ticket` and `new_ticket` fields in `kerberos.log` with `auth_ticket_sha256` and `new_ticket_sha256` via `protocols/krb/ticket-logging.zeek`; the `protocols/krb/md5-ticket-logging.zeek` policy script restores the original MD5 fields.

— Names old/new fields and both migration scripts. launch-20260917-0d171be8

04 **SSH host key fingerprint upgraded to SHA256** BREAKING

85

Replaces the `host_key` fingerprint field in `ssh.log` with a new `host_key_fingerprint` field using SHA256 in the same format as `ssh-keygen -l`; the `protocols/ssh/md5-host-key-logging.zeek` policy script restores the old field.

— Names old/new fields, output format, and restore script. launch-20260917-0d171be8

05 **zeek.conf cluster config with systemd orchestration** NEW

75

Adds a new declarative `zeek.conf` cluster configuration file consumed by Zeek's `systemd` generator to produce systemd unit files for orchestrating Zeek clusters, including multi-machine deployments.

— Names the config file and generator but gives no worked example.

launch-20260917-0d171be8

06 **ZeroMQ becomes default cluster communication backend** BREAKING

70

Zeek 9 switches the default cluster communication backend from Broker to ZeroMQ, enabling direct worker-to-worker eventing, cluster-wide event publishing, and single-point cluster event tapping; scripts or packages that rely on Broker-specific eventing must re-enable the Broker backend explicitly.

— Names backend switch and effect but no re-enable config key given.

launch-20260917-0d171be8

07 **DNS opcode fields for RFC 1996 notifications** NEW 70

Adds `opcode` and `opcode_name` fields to `dns.log` to support DNS notifications (RFC 1996); removable via the `protocols/dns/disable-opcode-log-fields.zeek`

policy script.

— Names fields, RFC reference, and removal script. launch-20260917-0d171be8

08 **Multicast logging improvements** NEW 60

Adds a new `multicast_participants.log` (via the `protocols/conn/multicast-participants.zeek` policy script) reporting multicast speakers based on IGMP activity, and logs multicast DNS transactions as separate transactions since they occur on separate network flows.

— Names log and policy script but no usage example. launch-20260917-0d171be8

09 **Spicy 1.17.0 upgrade with new optimizer** IMPROVED 60

Ships Spicy 1.17.0 with a control-flow-graph-based optimizer delivering 20–60% microbenchmark performance improvements and adds Windows (MSVC) build support.

— Gives concrete performance numbers but no upgrade steps. launch-20260917-0d171be8

THINNER COVERAGE BELOW

10 **`?as` type-conversion operator** NEW 55

Adds the `?as` operator to Zeek's scripting language for safely checking type-conversion compatibility at runtime.

— Names the operator and purpose but no code example. launch-20260917-0d171be8

11 **LDAP analyzer forwards SASL payloads** IMPROVED 50

Extends the LDAP analyzer to forward SASL payloads to GSSAPI and NTLM parsers, enhancing logging into `ntlm.log` and `kerberos.log`.

— Names target logs but no flag or config to enable. launch-20260917-0d171be8

12 **Native encapsulated-traffic ingestion for cloud** NEW 45

Ships a new packet source enabling Zeek clusters to act as a UDP server that ingests encapsulated packet streams directly for cloud environments, without packet-sniffing workarounds.

— Describes capability but gives no config key or command. launch-20260917-0d171be8

13 **NTP mode 6/7 logs** NEW 40

Adds new `ntp_control` and `ntp_private` logs for NTP modes 6 and 7 respectively.

— Names two new logs but no further mechanism given. launch-20260917-0d171be8

14 **ZIP file analyzer** NEW 40

Adds a new, optional ZIP file analyzer that reports contained files to `files.log`.

– Names the target log but no enable flag given. launch-20260917-0d171be8

15 **TCP support in Syslog analyzer** IMPROVED 25

Supports TCP in the Syslog analyzer.

– Bare statement with no mechanism or config given. launch-20260917-0d171be8

16 **FTP login/logout-only session logging** IMPROVED 25

Logs FTP sessions that feature only login and logout.

– Bare behavior note without mechanism given. launch-20260917-0d171be8

17 **JSON logging binary data configuration options** IMPROVED 25

Adds multiple configuration options for handling binary data in JSON logging, resolving a long-standing ambiguity.

– No option names given, only prose description. launch-20260917-0d171be8

└──▷ **BREAKING ON UPGRADE**

! The `host_key` field in `ssh.log` is replaced by `host_key_fingerprint` (SHA256, `ssh-keygen -l` format); any pipeline keying on `host_key` will break unless the `protocols/ssh/md5-host-key-logging.zeek` policy script is loaded to restore it.

! The `auth_ticket` and `new_ticket` fields (MD5 hashes) in `kerberos.log` are replaced by `auth_ticket_sha256` and `new_ticket_sha256` (SHA256) via `protocols/krb/ticket-logging.zeek`; load `protocols/krb/md5-ticket-logging.zeek` to restore the original fields.

! The default cluster communication backend is now ZeroMQ instead of Broker; scripts or packages that rely on Broker-specific eventing must re-enable the Broker backend explicitly.

announced · Sep 17, 2026 BLOG ↗

Zeek 9 ships ZeroMQ cluster comms by default, native encapsulated-traffic ingestion, systemd orchestration via zeek.conf, and new log fields.

└──▷ **TRY IT**

Enable ECH extension logging in ssl.log to capture HPKE KDF and AEAD identifiers from TLS 1.3 Encrypted Client Hello handshakes.

```
$ @load protocols/ssl/ssl-log-ext
```

Cap the software framework's detection cache to limit memory use and reduce noisy software-version log churn in high-volume environments.

/etc/zeek/local.zeek

```
redef Software::max_software_cache_size = 5000;
redef Software::parse_cache_interval = 1min;
redef Software::found_cache_interval = 5min;
```

Adds a new declarative `zeek.conf` cluster configuration file consumed by Zeek's `systemd` generator to produce systemd unit files for orchestrating Zeek clusters, including multi-machine deployments.

Adds `Software::parse_cache_interval`, `Software::found_cache_interval`, and `Software::max_software_cache_size` tunables to the software framework for finer control over version-detection caching.

Adds `ech_kdf_id` and `ech_aead_id` fields to `ssl.log` (via the `protocols/ssl/ssl-log-ext.zeek` policy script) reporting HPKE KDF and AEAD identifiers from the TLS 1.3 ECH extension.

Adds new `ntp_control` and `ntp_private` logs for NTP modes 6 and 7 respectively.

Adds a new `multicast_participants.log` (via the `protocols/conn/multicast-participants.zeek` policy script) reporting multicast speakers based on IGMP activity.

+14 more

└──▷ BREAKING ON UPGRADE

- **!** The `host_key` field in `ssh.log` is replaced by `host_key_fingerprint` (SHA256, `ssh-keygen -l` format); any pipeline keying on `host_key` will break unless the `protocols/ssh/md5-host-key-logging.zeek` policy script is loaded to restore it.
- **!** The `auth_ticket` and `new_ticket` fields (MD5 hashes) in `kerberos.log` are replaced by `auth_ticket_sha256` and `new_ticket_sha256` (SHA256) via `protocols/krb/ticket-logging.zeek`; load `protocols/krb/md5-ticket-logging.zeek` to restore the original fields.
- **!** The default cluster communication backend is now ZeroMQ instead of Broker; scripts or packages that rely on Broker-specific eventing must re-enable the Broker backend explicitly.

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Aviatrix

Feedback ⓘ 10.1.0 NOTES ↗ RANK

Welcome to the comprehensive documentation for Aviatrix's cloud networking platform. Here you'll find everything you need to deploy, configure, and manage your multi-cloud networking infrastructure.

> Aviatrix's 10.1.0 controller release is a major security and networking overhaul — FIPS 140-3 validated IPsec, end-to-end VRF multi-tenancy, DCF East-West L7 inspection, Premium IPS feeds and Flexible Gateway Deployment — bundled with dozens of Kubernetes, multi-cloud networking and CoPilot management updates, alongside a new Egress Security Score dashboard and Apple Silicon support in the VPN Client.

└──▷ **WHAT SHIPPED · 27 FEATURES** most completely described first ? what's the number?

01 **DCF intra-cluster enforcement for EKS** NEW 80

Adds DCF intra-cluster enforcement for EKS, enabled per-cluster from CoPilot under `Security > Distributed Cloud Firewall > Settings > Intra-Cluster`, supporting Calico, Cilium, Antrea, AWS VPC CNI (v1.14+), Weave Net, and Canal.

Enable DCF intra-cluster policy enforcement on an EKS cluster so that Distributed Cloud Firewall controls all east-west pod traffic via existing CNI NetworkPolicy support.

> 📍 In CoPilot, go to Security › Distributed Cloud Firewall › Settings › Intra-Cluster, select the target EKS cluster, and enable enforcement.

— Exact UI path, six named CNIs, and a step-by-step example 10.1.0

02 **Independent gateway instance sizing** NEW 80

Adds independent gateway instance sizing, allowing a single instance to be resized from CoPilot at `Cloud Fabric > Gateways > (transit or spoke) > gateway group > Instances` without affecting HA or the group default size; surfaces a Multiple Sizes label when a group has diverged.

Resize a single gateway instance without touching HA peers or changing the group default size, useful when one instance needs more capacity after a traffic spike.

> 📍 In CoPilot, go to Cloud Fabric › Gateways › (transit or spoke) › gateway group › Instances tab, select the individual instance, and choose a new size.

— Exact UI path plus a worked resize example 10.1.0

03 **Egress Security Score dashboard** NEW 75

Adds the `Security > Egress > Overview` tab (requires Controller 8.0 and DCF enabled) showing an Egress Security Score card for VPC/VNets and a breakdown of Unprotected, Monitored, Partially Protected, Protected, No Egress, and Unknown states. Adds an `Egress Security Score Explained` dialog (via `Learn More`) detailing how the score is calculated, and a `Manage VPC/VNets` shortcut that navigates directly to the Egress VPC/VNets tab for remediation.

Check your organisation's overall egress posture and identify unprotected VPC/VNets at a glance after upgrading to Controller 8.0 with DCF enabled.

> 📍 In the Aviatrix Controller, go to Security › Egress › Overview. Review the Egress Security Score card and the Protected VPC/VNets chart to see which VPCs/VNets are Unprotected, Monitored, Partially Protected, or Protected. Click `Manage VPC/VNets` to open the Egress VPC/VNets tab and begin remediating gaps.

— Exact UI path, states, and a worked usage example `product docs`

---

04 **DCF Spoke Local Breakout** NEW 75

Adds DCF Spoke Local Breakout, recording the egress path taken in every DCF log entry; requires Single IP SNAT (or customized SNAT), the 10.1 gateway image, and a supported instance size on Aviatrix Transit (FireNet) and AWS TGW centralized topologies.

— Precise prerequisites listed, but no enablement steps `10.1.0`

---

05 **DCF Quick Time-to-Value dashboard** NEW 75

Adds DCF Quick Time-to-Value dashboard with a system-defined, Aviatrix-managed read-only threat intelligence ruleset (defaults to Watch Mode) and ten Recommended Security Control cards covering ToR traffic, sanctioned-country traffic, cleartext legacy protocols, cryptomining pools, and outbound proxy/relay; requires Controller 8.2+ and CoPilot 4.37.

— Named cards and version requirements, no navigation path `10.1.0`

---

06 **FIPS build requires 2048-bit UserVPN CA key** BREAKING 70

The FIPS build of VPN Client 3.0.2 requires the VPN gateway's UserVPN CA to use an RSA key of at least 2048 bits; gateways on Controller versions earlier than 8.2, or Controller 8.2 whose UserVPN CA has not been rotated, may use a 1024-bit key (generated on Aviatrix 7.1 or earlier) that the FIPS 140-3 module will reject, preventing the client from verifying the gateway server certificate.

— Explains exact failure mode and versions affected `product docs`

---

07 **Premium IPS Rule Feeds (ET Pro)** NEW 70

Adds `Security > Distributed Cloud Firewall > IPS > Rule Feeds` (labeled ET Pro in CoPilot) as a Premium IPS feed applied alongside custom Suricata rulesets, with traffic actions governed by severity thresholds in the active IPS Profile.

Activate the Premium ET Pro IPS rule feed so that matched traffic is handled according to your active IPS Profile severity thresholds, without replacing any custom Suricata rules you have uploaded.

> 📍 In CoPilot, go to Security › Distributed Cloud Firewall › IPS › Rule Feeds; the ET Pro feed is listed as the default entry — confirm it is enabled, then verify or adjust severity thresholds in your active IPS Profile.

— Named endpoint path and worked example, but no threshold defaults `10.1.0`

---

08 **VRF for end-to-end multi-tenancy** NEW 70

Adds VRF for multi-tenancy with end-to-end propagation across transit-to-transit peering, BGP-spoke-to-transit peering, and intra-group backup links; enabled by a Controller feature flag (off by default) plus per-attachment opt-in, supporting up to 200 network domains.

— Mechanism and 200-domain limit given, but flag name not specified `10.1.0`

---

09 **DCF for Private Kubernetes Clusters (Preview)** NEW 70

Adds DCF for Private Kubernetes Clusters (Preview), enabled per cluster in CoPilot or Terraform, supporting EKS, AKS, and GKE; requires an Aviatrix spoke gateway in the same VPC with a security-group rule allowing TCP port 443 to the Kubernetes API server.

— Exact prerequisite (TCP 443 rule) named, no walkthrough `10.1.0`

---

10 **Flexible Gateway Deployment** NEW 60

Adds Flexible Gateway Deployment (private subnet, no public IP), enabled per gateway group at creation time on AWS (including GovCloud) and Azure, via CoPilot and Terraform.

— Clouds and tools named, no exact config keys `10.1.0`

---

11 **DCF East-West Layer 7 inspection for AWS** NEW 60

Adds DCF East-West Layer 7 inspection for AWS, with symmetric routing enabled per spoke group in CoPilot or Terraform, also reducing cross-AZ data transfer cost and latency.

— Names benefit and tools but not exact config field `10.1.0`

---

12 **GCP Global VPC Regional Spoke Egress (Phase 2)** IMPROVED

60

Adds GCP Global VPC Regional Spoke Egress (Phase 2), letting operators choose which regions egress locally vs. through FireNet, with both modes coexisting in the same Global VPC; the CoPilot Egress page now shows Global VPCs in a nested tree displaying each region's egress mode.

— UI location named, no exact setting path `10.1.0`

THINNER COVERAGE BELOW

13 **SNAT and HPE support on Azure spokes** NEW

55

Adds SNAT and HPE support on Azure spokes with subnet groups for both HPE and non-HPE deployments, configurable via the existing Controller API/CLI or Terraform.

— Names config surfaces but no specific fields `10.1.0`

14 **Smart Gateways Phase 3 (Preview)** NEW

55

Adds Smart Gateways Phase 3 (Preview) covering transit gateways, Edge as Spoke, Edge as Transit, and BGP spoke gateways, controlled by three tiered feature flags (fabric-wide kill switch, per-group underlay flag, per-group route-resolver flag), all off by default.

— Flags named but no activation procedure `10.1.0`

15 **VPN Client 3.0.2 Apple Silicon and Ubuntu 26.04 support** NEW

50

Packages the macOS installer as a Universal binary that runs natively on both Intel and Apple Silicon (ARM64) Macs without Rosetta 2, and adds support for Ubuntu 26.04 alongside Ubuntu 22.04 and 24.04.

— Names platforms covered but no install/upgrade command given `product docs`

16 **FIPS 140-3 validated IPsec data plane** NEW

50

Adds a FIPS 140-3 validated IPsec data plane as a separate FIPS-validated gateway build; existing gateways must be replaced (not toggled) to enter FIPS mode.

— States migration requirement but no replacement procedure `10.1.0`

17 **Kubernetes Discovery and SmartGroups reach GA** IMPROVED

45

Promotes Kubernetes Discovery and SmartGroups to general availability, enabling DCF rules to enforce policy across both Kubernetes Pods/Services and VM workloads.

— Describes scope of GA but no setup steps `10.1.0`

18 **AZ Affinity Across Transit** NEW

45

Adds AZ Affinity Across Transit, enabled per peering (disabled by default) on AWS same-region HPE peerings.

— Scope named but no enablement steps 10.1.0

19 **AWS inter-region jumbo frame support** NEW 45

Adds opt-in AWS inter-region jumbo frame support for transit peering (AWS only, excluding partitions such as AWS GovCloud that do not support jumbo frames).

— Names exclusions but no config step 10.1.0

20 **AKS onboarding with Microsoft Entra ID** NEW 45

Adds AKS onboarding with Microsoft Entra ID, with the Controller auto-detecting the Entra ID setting on the AKS cluster.

— Describes auto-detection but no onboarding steps 10.1.0

21 **Private Mode removed, must migrate to Flexible Gateway** 45

BREAKING

Private Mode is fully removed; deployments still on Private Mode cannot upgrade in place and must be rebuilt as Flexible Gateway deployments.

— Clear migration requirement, no rebuild procedure detail 10.1.0

22 **Make Before Break Gateway Replacement (Preview)** NEW 40

Adds Make Before Break Gateway Replacement (Preview), preserving the gateway's external IP through cut-over so partners and firewalls need no reconfiguration, with automatic eligibility checking per gateway.

— Behaviour described but no eligibility criteria or steps 10.1.0

23 **Controller UI features move exclusively to CoPilot** BREAKING 35

Controller UI features formerly managed in the Controller UI are now managed exclusively from CoPilot; onboarding flows for new and existing deployments now route through CoPilot.

— States the shift but not which specific features moved 10.1.0

24 **Azure v6 instance support** NEW 30

Adds Azure v6 instance support for gateways on Controller 10.1.0 and later.

— Names instance family and version, nothing else 10.1.0

25 **Ubuntu 20.04 support deprecated** DEPRECATED 25

Ubuntu 20.04 support is deprecated in VPN Client 3.0.2.

— Bare deprecation notice, no timeline or migration path `product docs`

26 **DCF CSP Exclusion List reaches GA** IMPROVED 25

Promotes DCF CSP Exclusion List to general availability.

— Bare GA announcement with no mechanism `10.1.0`

27 **Universal Installer for Aviatrix OpenVPN Client** NEW 25

Adds Universal Installer for the Aviatrix OpenVPN Client.

— One-line mention with no further detail `10.1.0`

└──▷ BREAKING ON UPGRADE

! The FIPS build of VPN Client 3.0.2 requires the VPN gateway's UserVPN CA to use an RSA key of at least 2048 bits; gateways on Controller versions earlier than 8.2, or Controller 8.2 whose UserVPN CA has not been rotated, may use a 1024-bit key (generated on Aviatrix 7.1 or earlier) that the FIPS 140-3 module will reject, preventing the client from verifying the gateway server certificate.

! Ubuntu 20.04 support is deprecated in VPN Client 3.0.2.

! Private Mode is fully removed; deployments still on Private Mode cannot upgrade in place and must be rebuilt as Flexible Gateway deployments.

! Controller UI features formerly managed in the Controller UI are now managed exclusively from CoPilot; onboarding flows for new and existing deployments now route through CoPilot.

`10.1.0` NOTES ↗ LATEST STABLE

Aviatrix 10.1.0 adds FIPS 140-3 IPsec, end-to-end VRF multi-tenancy, DCF East-West L7 inspection, Premium IPS feeds, and Flexible Gateway Deployment.

└──▷ HOW TO FIND IT

Enable DCF intra-cluster policy enforcement on an EKS cluster so that Distributed Cloud Firewall controls all east-west pod traffic via existing CNI NetworkPolicy support.

📍 In CoPilot, go to Security › Distributed Cloud Firewall › Settings › Intra-Cluster, select the target EKS cluster, and enable enforcement.

Activate the Premium ET Pro IPS rule feed so that matched traffic is handled according to your active IPS Profile severity thresholds, without replacing any custom Suricata rules you have uploaded.

📍 In CoPilot, go to Security › Distributed Cloud Firewall › IPS › Rule Feeds; the ET Pro feed is listed as the default entry — confirm it is enabled, then verify or adjust severity thresholds in your active IPS Profile.

Resize a single gateway instance without touching HA peers or changing the group default size, useful when one instance needs more capacity after a traffic spike.

📍 In CoPilot, go to Cloud Fabric › Gateways › (transit or spoke) › gateway group › Instances tab, select the individual instance, and choose a new size.

Adds `Security > Distributed Cloud Firewall > IPS > Rule Feeds` (labeled `ET Pro` in CoPilot) as a Premium IPS feed applied alongside custom Suricata rulesets, with traffic actions governed by severity thresholds in the active IPS Profile.

Adds DCF intra-cluster enforcement for EKS, enabled per-cluster from CoPilot under `Security > Distributed Cloud Firewall > Settings > Intra-Cluster`, supporting Calico, Cilium, Antrea, AWS VPC CNI (v1.14+), Weave Net, and Canal.

Adds VRF for multi-tenancy (end-to-end propagation) across transit-to-transit peering, BGP-spoke-to-transit peering, and intra-group backup links; enabled by a Controller feature flag (off by default) plus per-attachment opt-in, supporting up to 200 network domains.

Adds Flexible Gateway Deployment (private subnet, no public IP), enabled per gateway group at creation time on AWS (including GovCloud) and Azure, via CoPilot and Terraform.

Adds FIPS 140-3 validated IPsec data plane as a separate FIPS-validated gateway build; existing gateways must be replaced (not toggled) to enter FIPS mode.

+16 more

└──▷ **BREAKING ON UPGRADE**

! Private Mode is fully removed; deployments still on Private Mode cannot upgrade in place and must be rebuilt as Flexible Gateway deployments.

! Controller UI features formerly managed in the Controller UI are now managed exclusively from CoPilot; onboarding flows for new and existing deployments now route through CoPilot.

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ Fastly Terraform Provider 9.7.0 (2026-09-17)

BLOG ↗ RANK

Fastly's edge cloud platform delivers faster, safer, and more scalable sites and apps to customers. Elevate your edge CDN, video delivery, security, and more.

> Fastly's window centered on expanded DDoS protection observability and Terraform provider updates, adding ten new challenge metrics to its analytics APIs and new Terraform resources for AI Runtime Control and DDoS client challenges.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 **DDoS protection challenge metrics in analytics APIs** NEW

60

Adds ten new DDoS protection challenge metrics to the Historical Stats API and Real-Time Analytics API, covering challenge and token creation, validation outcomes, and Private Access Token (PAT) challenges issued during DDoS protection rule enforcement.

– Names two APIs and metric categories but not exact metric names

Add DDoS protection challenge metrics t…

THINNER COVERAGE BELOW

02 **New Terraform provider resources for AI Runtime Control and DDoS challenge mode**

NEW

55

Fastly Terraform Provider 9.7.0 adds AI Runtime Control resources and support for `client_challenge` mode in the `product_enablement/ddos_protection` resource.

– Names exact resource and mode flag but little mechanism detail

Fastly Terraform Provider 9.7.0 (2026-0…

> Fastly Terraform Provider 9.7.0 (2026-09-17) BLOG ↗ LATEST STABLE
>
> Fastly Terraform Provider 9.7.0 adds AI Runtime Control resources and `client_challenge` DDoS protection mode.

WAS THIS USEFUL? 👍 👎

Rspamd

Feedback ⓘ 4.2.0 NOTES ↗ RANK

Rapid spam filtering system.

> Rspamd 4.2.0 replaces its per-class symbol schedulers with a single dependency-ordered execution plan and rewrites TLD/public-suffix resolution away from hyperscan matching, while adding bounded content extraction for SVG, OOXML and DOCX attachments, a reworked fuzzy-storage counting and health-monitoring pipeline, and hard limits on parsing untrusted UCL/JSON/msgpack input.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 **Fuzzy storage counting, tracking and health monitoring**

IMPROVED 92

Adds a `count_scan` section to the fuzzy worker config with tunable `pass interval` (default 4 hours), `SCAN COUNT` (default 1000), and `duty cycle` (default 10%), replacing the broken `<prefix>_count` Redis counter with a periodic SCAN-based keyspace walk; sampled statistics from this are configured via `stats_sample` and reported under `<prefix>_stats`, `/fuzzystat`, and `rspamadm fuzzystat`. Per-source IP tracking is restored via worker-level and per-key `max_ips_per_key` / `max_ips` (now lazy and off by default for the default key), with `ips_inserted` and `ips_overflow` fields, and a table exceeding eight times its eviction cap in an hour is dropped and flagged. Fuzzy storage key policy now applies after decryption so a valid customer key is no longer dropped by an IP ban, and health monitoring is added via `fuzzy_check.ping_storage_all` and the `/plugins/fuzzy/status` endpoint, with per-server liveness marks, down badges, and an on-demand Check button in the WebUI.

Restore per-IP fuzzy storage tracking for the default key and cap it at 1024 entries after the lazy-statistics change makes it opt-in.

```ini
max_ips_per_key = 1024;
```

Query the live health of all configured fuzzy storages from the command line to verify reachability before a maintenance window.

```
$ curl -s http://localhost:11334/plugins/fuzzy/status | python3 -m json.tool
```

Pull sampled fuzzy storage statistics — including per-source IP counts — to audit hash-count accuracy after migrating from the old counter-based approach.

```
$ rspamadm fuzzystat
```

Cap per-key source IP tracking in the fuzzy storage to limit memory use on high-traffic deployments.

```ucl
max_ips_per_key = 256;

key {
  id = "<your-key-id>";
  max_ips = 64;
}
```

— Names every config key, endpoint, and gives runnable commands. 4.2.0

02 **Public suffix and TLD resolution rewrite** BREAKING 90

Replaces the old multipattern/hyperscan TLD lookup with a longest-suffix hostname matcher against the effective TLDs file, honouring `!` exception rules that were previously skipped. Exposed to Lua as rspamd_tld_lookup(), url:get_public_suffix() with match flags, and the `get_public_suffix` selector transform, supporting wildcards, exceptions, and ICANN/private splits. Two-pass URL discovery in free text now uses a static matcher plus dot-anchored suffix probing, eliminating the old 10,500-pattern hyperscan database (previously 1.3s to compile, 6.5MiB serialised); RBL URL composition maps are now matched via `rspamd_tld_lookup` instead of per-TLD regexp buckets. `filter = "top"` in multimap now returns the full public suffix (e.g. `co.uk`, `com.au`) instead of only the last label, and a host that is itself a public suffix resolves to the whole host — map files must be audited and updated before upgrading.

Audit multimap rules that use `filter = 'top'` before upgrading, to find map files that need bare multi-label labels replaced with full suffixes.

```
$ grep -rn 'filter *= *"top"' /etc/rspamd/local.d/
/etc/rspamd/override.d/
```

Verify that TLD resolution returns the correct eSLD for domains under PSL exception rules after the upgrade.

```
$ rspamadm lua -e
'print(require("rspamd_util").get_tld("www.city.kobe.jp"))'
```

— Names Lua API, config impact, and gives audit/verify commands. `4.2.0`

---

03 **Single symbol execution scheduler** BREAKING 85

Replaces per-class symbol schedulers with one execution plan computed at cache initialisation, viewable via `rspamadm configdump -d` and `rspamadm configdump -e`, and introducing `suppressed` and `skipped` item states. Prefilters can now legitimately depend on filters, and invalid register_dependency() edges (e.g. a connfilter depending on a filter in the wrong stage) are rejected at configuration time instead of silently accepted; `ignore_passthrough` filters now apply correctly after a pre-result so task:set_pre_result() no longer skips the filters stage.

Inspect the new unified execution plan to verify symbol ordering and stage assignments after upgrading or adding custom rules.

```
$ rspamadm configdump -e
```

— Names commands, new states, and dependency-validation mechanism. `product docs`

---

04 **Bounded content extraction for SVG, OOXML and DOCX** NEW 73

Adds SVG attachment content extraction with smuggling-indicator symbols `SVG_CONTENT`, `SVG_SCRIPT`, `SVG_FOREIGN_OBJECT`, `SVG_DATA_URI`, `SVG_FORM`, controlled via a new `lua_content { svg { ... } }` config block for per-task budget and timeout. Adds XLSX/PPTX extraction in the bounded OOXML pipeline covering shared strings, worksheets, Excel 4.0 macrosheets, DrawingML, slide notes, and auto-executing `Auto_Open` defined names, surfaced as `OOXML_MACROS`, `OOXML_OLE_OBJECT`, `OOXML_REMOTE_TEMPLATE`, `OOXML_EXTERNAL_DATA`. Adds bounded DOCX extraction with the core rewritten in C++, flagging unprocessable attachments instead of silently ignoring them.

— Names every new symbol and the SVG config block. `4.2.0`

---

05 **Bounded parsing of untrusted UCL/JSON/msgpack input** BREAKING 70

Adds ucl.untrusted_parser(), parser:set_limits(), and parser:get_limits() Lua APIs for bounded parsing of untrusted network replies, covering `checkv3`, controller, and proxy body parsers. Enforces depth 64, 1 million elements, 1KB keys, 16MB strings

(individually overridable via `max_string_length`), and min(96*inlen+256KB, 64MB) of tree; requests exceeding these limits are now rejected rather than accepted.

— Names the API calls and exact numeric limits. 4.2.0

THINNER COVERAGE BELOW

06 **SPF permerror enforcement for duplicate records** BREAKING

54

Enforces RFC 7208 §4.5: domains publishing more than one `v=spf1` record now return `permerror` instead of being evaluated on whichever record the resolver answered first, changing DMARC outcomes for affected mail. `V=SPF1` now counts as an SPF record (case-insensitive per RFC), while `v=spf10` no longer does.

— Names the RFC section and exact matching rule change. product docs

07 **WebUI status and Bayes panel improvements** IMPROVED

51

The Status servers table gains three-state health, per-server details rows, load and latency columns, version and config drift indicators, a Writable badge, and an `All SERVERS` cluster row. The Bayes panel gains a learns balance bar and a `min_learns` status badge.

— Names specific UI elements but only a navigation path. 4.2.0

08 **DKIM alignment scoring** BREAKING

50

Adds a dedicated `R_DKIM_ALIGNED` symbol (weight -0.1) that separately scores DKIM signatures aligned with the author domain, decided in the DKIM module; `R_DKIM_ALLOW` weight is reduced from -0.2 to -0.1, so installations that tuned it must recalibrate both symbols together.

— Names both symbols and exact weight change. 4.2.0

09 **ClickHouse schema upgrade to version 12** IMPROVED

48

Upgrades the ClickHouse schema to version 12, adding bloom-filter data-skipping indexes (granularity 4) on `MessageId`, `From`, `MimeFrom`, `IP`, `Urls.Tld`, `Attachments.Digest`, and `Subject` for faster point lookups.

— Names schema version and all indexed columns. 4.2.0

10 **HTML part injection from Lua** IMPROVED

44

Extends task:inject_part() to accept `'html'` as a part type, routing injected content through the full HTML parser so it gets URL, feature, and CTA-link extraction like native HTML parts.

— Names the exact API call and new part-type value. 4.2.0

11 **CSS compound selector and combinator support** NEW 43

Adds CSS compound selector and combinator evaluation, supporting `div.mainbox`, `div p`, `div > p`, `h1 + p`, and `h1 ~ p` patterns.

— Lists exact selector patterns but no usage context. 4.2.0

12 **redis_list pusher for metadata_exporter** NEW 30

Adds `redis_list` pusher support to `metadata_exporter`.

— Single named surface, no further mechanism described. 4.2.0

13 **lua_aliases mailbox identity classes** NEW 27

Adds `lua_aliases` mailbox identity and equivalent domain classes for comparing addresses.

— Bare name of module, no mechanism or usage given. 4.2.0

14 **Upstream pinning for lua_redis** IMPROVED 23

Adds upstream pinning for `lua_redis` requests.

— One-line mention with no mechanism or scope. 4.2.0

└──▷ BREAKING ON UPGRADE

! The single symbol scheduler replaces per-class schedulers; symbol dependency ordering and stage assignment change — a prefilter that previously pulled a filter into the prefilter stage now correctly executes after it, which may alter scan results on custom rule sets.

! Public suffix lookup now honours `!` exception rules that were previously skipped; a host that is itself a public suffix now resolves to the whole host rather than a shorter suffix, changing URL and RBL query behaviour.

! Two-pass URL discovery replaces the hyperscan-based TLD multipattern database; URL detection results may differ from previous versions.

! The `multimap` `top` filter now returns the full public suffix (e.g. `com.au`) instead of only the last label (`au`); maps using this filter must be updated to list full suffixes before upgrading.

! RBL URL composition maps are now matched by the new `rspamd_tld_lookup` rule set instead of per-TLD regexp buckets; wildcard and exception handling changes as described.

! Fuzzy Redis hash counts are no longer maintained by the update script; the `<prefix>_count` key is no longer incremented/decremented on add/delete. Negative counters from older versions are reported as 0 until the first SCAN pass completes.

! `R_DKIM_ALLOW` weight changes from -0.2 to -0.1; the new `R_DKIM_ALIGNED` symbol carries the other -0.1. Installations that tuned `R_DKIM_ALLOW` must revisit both symbols together.

! `filter = "top"` in multimap now returns the full public suffix (e.g. `co.uk`) instead of the last label (`uk`); map files listing bare multi-label labels like `au`, `uk`, `br`, `jp` must be updated to the full suffixes they intend.

- ! Domains publishing more than one `v=spf1` record now receive `permerror` instead of a pass/fail; such mail also changes DMARC outcome. Audit and merge duplicate SPF records before upgrading.
- ! `V=SPF1` now counts as an SPF record (case-insensitive match per RFC), and `v=spf10` no longer counts as one.
- ! `R_DKIM_ALLOW` score changes from -0.2 to -0.1; rules or thresholds tuned around the old score must be recalibrated (add `R_DKIM_ALIGNED` at -0.1 to restore the original total for aligned mail).
- ! Invalid rspamd_config:register_dependency() edges (e.g. a connfilter or prefilter depending on a filter in the wrong stage) are now rejected at configuration time and dropped, rather than being silently accepted.
- ! Fuzzy hash count resets to 0 after upgrade and remains 0 until the first `count_scan` SCAN pass completes (up to 4 hours by default); monitoring thresholds on `<prefix>_count` must be re-baselined.
- ! Per-source fuzzy statistics are now off by default for the default key; operators relying on per-IP tracking for the default key must set `max_ips_per_key` and add an explicit `max_ips` extension in the keys file.
- ! UCL untrusted parsing now rejects request bodies with strings exceeding 16 MB or nesting deeper than 64 levels; `checkv3`, controller, and proxy requests previously accepted with such payloads are now rejected.
- ! Fuzzy storage IP bans no longer block clients holding a valid customer key; operators who used IP bans to cut off such clients must revoke or expire the key itself instead.
- ! TLD resolution now honours PSL `!` exception rules and resolves a host that is itself a public suffix to the whole host; domain allowlist/blocklist entries keyed on eSLD under an exception rule may match a different key after upgrade.
- ! Maps used with the `top` filter in multimap must now list full suffixes (e.g. `com.au` instead of `au`) — the old TLD table is replaced by the public suffix lookup.
- ! The 10.5k-pattern TLD hyperscan database is removed from the scanner; free-text URL discovery now runs two passes (static matcher plus dot-anchored suffix labels) via the public suffix list.
- ! RBL URL composition maps and host TLD resolution now go through the public suffix lookup; `!` exception rules are honoured and a host that is itself a public suffix resolves to the whole host.
- ! Fuzzy Redis hash counts are now derived from a periodic SCAN instead of the counter maintained by the update script; negative counters left by older versions are reported as 0.
- ! `R_DKIM_ALLOW` weight changes from -0.2 to -0.1; alignment credit moves to the new `R_DKIM_ALIGNED` symbol.

4.2.0 NOTES ↗ LATEST STABLE

Rspamd 4.2.0 adds SVG/OOXML content extraction, public suffix Lua API, fuzzy storage health UI, and untrusted UCL parsing limits.

└──▷ GET THIS VERSION

```
$ git clone --branch 4.2.0 https://github.com/rspamd/rspamd.git
# already have the repo? check out this version:
$ git checkout 4.2.0
```

└──▷ **TRY IT**

Query the live health of all configured fuzzy storages from the command line to verify reachability before a maintenance window.

```
$ curl -s http://localhost:11334/plugins/fuzzy/status | python3 -m
json.tool
```

Pull sampled fuzzy storage statistics — including per-source IP counts — to audit hash-count accuracy after migrating from the old counter-based approach.

```
$ rspamadm fuzzystat
```

Cap per-key source IP tracking in the fuzzy storage to limit memory use on high-traffic deployments.

```ucl
max_ips_per_key = 256;

key {
  id = "<your-key-id>";
  max_ips = 64;
}
```

Exposes public suffix list semantics to Lua as rspamd_tld_lookup(), url:get_public_suffix() with match flags, and the `get_public_suffix` selector transform, supporting wildcards, exceptions, and ICANN/private splits.

Adds ucl.untrusted_parser(), parser:set_limits(), and parser:get_limits() Lua API calls for bounded parsing of untrusted UCL/JSON/msgpack network input (controller and proxy bodies, checkv3 metadata).

Adds fuzzy storage health monitoring via `fuzzy_check.ping_storage_all` and the `/plugins/fuzzy/status` API endpoint, with per-server liveness marks, down badges, and an on-demand Check button in the WebUI.

Reports sampled fuzzy storage statistics from count scans (`count_scan.stats_sample`) in `/fuzzystat` and via `rspamadm fuzzystat`, with configurable per-source IP caps (`max_ips_per_key`, per-key `max_ips`).

Extracts SVG attachment content and smuggling indicators as new symbols: `SVG_CONTENT`, `SVG_SCRIPT`, `SVG_FOREIGN_OBJECT`, `SVG_DATA_URI`, `SVG_FORM`.

`+10 more`

└──▷ **BREAKING ON UPGRADE**

- ! Maps used with the `top` filter in multimap must now list full suffixes (e.g. `com.au` instead of `au`) — the old TLD table is replaced by the public suffix lookup.
- ! The 10.5k-pattern TLD hyperscan database is removed from the scanner; free-text URL discovery now runs two passes (static matcher plus dot-anchored suffix labels) via the public suffix list.
- ! RBL URL composition maps and host TLD resolution now go through the public suffix lookup; `!` exception rules are honoured and a host that is itself a public suffix resolves to the whole host.
- ! Fuzzy Redis hash counts are now derived from a periodic SCAN instead of the counter maintained by the update script; negative counters left by older versions are reported as 0.
- ! `R_DKIM_ALLOW` weight changes from -0.2 to -0.1; alignment credit moves to the new `R_DKIM_ALIGNED` symbol.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Anchore Grype

Feedback ⓘ v0.119.0 NOTES ↗ CODE ↗ RANK

A vulnerability scanner for container images and filesystems

> Grype added a config option to reveal vulnerability matches that its matchers had internally suppressed, such as distro-fixed and known false-positive records.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **Visibility into matcher-suppressed matches** NEW 86

A new `include-matcher-suppressions` config key (yaml: `include-matcher-suppressions`) exposes matches that matchers internally dropped — including distro-fixed and built-in false-positive records — in the `ignoredMatches` output, so they become visible when using `--show-suppressed`.

Permanently enable suppressed-matcher match visibility in your Grype config so every scan includes distro-fixed dropped matches in its ignored output.

yaml

```yaml
include-matcher-suppressions: true
show-suppressed: true
```

— Names exact config key, flag, and output field with usage example. v0.119.0

v0.119.0 NOTES ↗ CODE ↗ LATEST STABLE

Grype now surfaces distro-fixed suppressed matches via `--show-suppressed` using the new `include-matcher-suppressions` option.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.119.0 https://github.com/anchore/grype.git
# already have the repo? check out this version:
$ git checkout v0.119.0
```

└──▷ **USE IT**

Permanently enable suppressed-matcher match visibility in your Grype config so every scan includes distro-fixed dropped matches in its ignored output.

.grype.yaml

```
include-matcher-suppressions: true
show-suppressed: true
```

Adds `include-matcher-suppressions` config key (yaml: `include-matcher-suppressions`) to expose matches internally dropped by matchers — such as distro-fixed and built-in false-positive records — in the `ignoredMatches` output, making them visible when `--show-suppressed` is used.

WAS THIS USEFUL? 👍 👎

Anchore syft

Feedback ⓘ v1.52.0 NOTES ↗ RANK

CLI tool and library for generating a Software Bill of Materials from container images and filesystems

> Syft v1.52.0 adds a passthrough-exact-format output option that returns matching syft-json input unchanged, avoiding unnecessary re-encoding.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **passthrough-exact-format output format** NEW 70

A new `passthrough-exact-format` output format copies a matching `syft-json` input document to output byte-for-byte without re-encoding.

— Names exact flag and behavior but no usage example given v1.52.0

v1.52.0 NOTES ↗ LATEST STABLE

Syft v1.52.0 adds passthrough-exact-format to copy matching syft-json input to output unchanged.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v1.52.0 https://github.com/anchore/syft.git
# already have the repo? check out this version:
$ git checkout v1.52.0
```

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Wiz Defend

Feedback ⓘ changelog · Sep 15, 2026 BLOG ↗ RANK

Wiz Defend is a cloud security platform that identifies and remediates vulnerabilities and misconfigurations across cloud environments.

> Wiz Defend added a native integration with Google Security Operations, letting analysts install a one-click Content Pack that brings Wiz's AI threat analysis and case data directly into Google's SOC workflow.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **Google Security Operations integration via Wiz Content Pack**

NEW 73

Wiz Defend now ships a single-click install Content Pack for Google Security Operations that provides out-of-the-box rules, dashboards, search queries, playbooks, and response policies. Playbooks automatically fetch and display Wiz Blue Agent AI-powered threat analysis directly inside Google Security Operations cases, and a bidirectional real-time sync keeps status, severity, and comments consistent between Wiz Defend threats and Google Security Operations cases. A shared data model across both platforms lets analysts investigate Wiz threats directly within Google Security Operations.

— Explains mechanism and named components, but no config keys or API endpoints given.

2026-09-15-39a74171

changelog · Sep 15, 2026 BLOG ↗

Wiz Defend integrates with Google Security Operations via a Content Pack, bringing Blue Agent analysis, shared data model, and bidirectional sync.

Adds the Wiz Content Pack for Google Security Operations, providing out-of-the-box rules, dashboards, search queries, playbooks, and response policies in a single-click install.

Adds playbooks in Google Security Operations to automatically fetch and display Wiz Blue Agent AI-powered threat analysis directly in Google Security Operations cases.

Adds bidirectional real-time sync of status, severity, and comments between Wiz Defend threats and Google Security Operations cases.

Enables analysts to investigate Wiz threats directly within Google Security Operations via a shared data model across both platforms.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Threat Intelligence

Feedly Threat Intelligence

Feedback ⓘ

Pivot from an IP IoC to live host data in Censys (2026-09-17) BLOG ↗ RANK

Collect, analyze, share actionable intel faster - Feedly TI

> Feedly Threat Intelligence added a one-click pivot from IP indicators of compromise directly into live Censys host data, without requiring a separate Censys account.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **Censys pivot from IP IoCs** NEW 85

An Open in Censys button on IoC Insights Cards lets analysts jump from a flagged IP directly to its live Censys host record, showing open ports, running services, TLS certificates, ASN, hosting provider, and a CVEs tab with CVSS scores and KEV membership — no Censys account required. The same Open in Censys pivot was extended to the IoC table on Threat Actor and Malware Insights Cards, so analysts can pivot from any IP row across a full IoC list without opening each card individually, including a Network and Routing panel with ASN, Organization, Routing Prefix, and WHOIS data.

Verify whether a flagged IP is still live and enumerate its exposed services before blocking, to avoid false positives against reassigned addresses.

📍 In an IoC Insights Card for an IP indicator, click Open in Censys to view the host record — including open ports, running services, TLS certificates, ASN, and the CVEs tab with CVSS scores and KEV membership.

Pivot across all IoCs attributed to a threat actor at once to identify shared infrastructure without drilling into each indicator card individually.

📍 Open a Threat Actor Insights Card, scroll to the IoC table, and click Open in Censys on any IP row to pivot directly to that host's Censys record — including the Network and Routing panel for ASN, Organization, Routing Prefix, and WHOIS data.

Validate whether a reported IP is still live and check for shared infrastructure before actioning a block.

📍 In Feedly Threat Intelligence, open an IoC Insights Card for any IP indicator, then click Open in Censys to view the live host record, open ports, TLS certificates,

and certificate-sharing peers in Censys — no Censys account required.

Pivot across all IoCs attributed to a threat actor without opening individual cards, useful when triaging a batch of infrastructure indicators.

📍 In Feedly Threat Intelligence, open a Threat Actor Insights Card, scroll to the IoC table, and click `Open in Censys` on any IP row to jump directly to that host's Censys record.

— Names exact UI element, data fields, and navigation path across card types.

`Pivot from an IP IoC to live host data …`

`Pivot from an IP IoC to live host data in Censys (2026-09-17)` `BLOG ↗`

LATEST `STABLE`

Feedly IoC Insights Cards now link directly to Censys host data for one-click IP pivot without leaving the platform.

└──▷ **HOW TO FIND IT**

Validate whether a reported IP is still live and check for shared infrastructure before actioning a block.

📍 In Feedly Threat Intelligence, open an IoC Insights Card for any IP indicator, then click `Open in Censys` to view the live host record, open ports, TLS certificates, and certificate-sharing peers in Censys — no Censys account required.

Pivot across all IoCs attributed to a threat actor without opening individual cards, useful when triaging a batch of infrastructure indicators.

📍 In Feedly Threat Intelligence, open a Threat Actor Insights Card, scroll to the IoC table, and click `Open in Censys` on any IP row to jump directly to that host's Censys record.

Adds an `Open in Censys` button to IoC Insights Cards, enabling one-click pivot from a flagged IP to live Censys host data including open ports, running services, TLS certificates, ASN, and hosting provider.

Extends the `Open in Censys` pivot to the IoC table on Threat Actor and Malware Insights Cards, so analysts can pivot from a full IoC list without opening each individual card.

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 7.260917.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI 7.260917.0 adds multi-source OpenAEV security coverage analysis, SSVC-based vulnerability risk prioritization, a full user-merge framework, and simplified direct-to-PDF Fintel export, alongside new lock-contention telemetry and workflow publish actions.

└──▷ **WHAT SHIPPED · 6 FEATURES** most completely described first ? what's the number?

01 **User-merge framework with dry-run and RBAC merging** NEW 75

Introduces a user-merge framework supporting `dry-run` / `run` modes, a handler registry, scalar user reference rewriting, RBAC rights merging, and union/strict merge options.

– Names modes and mechanisms but no exact command syntax 7.260917.0

THINNER COVERAGE BELOW

02 **Direct HTML-to-PDF Fintel export** IMPROVED 50

Simplifies Fintel export by generating PDF directly from HTML without an intermediate step, streamlining intelligence report delivery.

– Explains before/after but no endpoint or UI path named 7.260917.0

03 **SSVC risk-based prioritization for vulnerabilities** NEW 45

Lands SSVC risk-based prioritization metrics for vulnerabilities behind a feature flag, introducing exploit-driven risk scoring to the vulnerability workflow.

– Feature-flagged; no flag name or scoring detail given 7.260917.0

04 **Workflow publish actions for Status objects** NEW 45

Adds workflow publish actions that create and resolve Status objects on state transitions.

– Names Status objects but no config or UI path given 7.260917.0

05 **Multi-OpenAEV security coverage results** NEW 40

Adds multi-OpenAEV security coverage results, enabling security coverage analysis across multiple OpenAEV sources simultaneously.

– Describes capability but no interface or config detail 7.260917.0

06 **Lock contention telemetry metrics** NEW 35

Exposes lock acquisition wait and contention metrics via telemetry to surface backend lock pressure.

— No metric names or dashboard location provided 7.260917.0

7.260917.0 NOTES ↗ LATEST STABLE

OpenCTI 7.260917.0 adds multi-OpenAEV security coverage, SSVC vulnerability prioritization, user-merge framework, and direct Fintel PDF export.

└──▷ GET THIS VERSION

```
$ git clone --branch 7.260917.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260917.0
```

Adds multi-OpenAEV security coverage results, enabling security coverage analysis across multiple OpenAEV sources simultaneously.

Lands SSVC risk-based prioritization metrics for vulnerabilities behind a feature flag, introducing exploit-driven risk scoring to the vulnerability workflow.

Introduces a user-merge framework with dry-run/run modes, a handler registry, scalar user reference rewriting, RBAC rights merging, and union/strict merge options.

Simplifies Fintel export by generating PDF directly from HTML without an intermediate step, streamlining intelligence report delivery.

Exposes lock acquisition wait and contention metrics via telemetry to surface backend lock pressure.

+1 more

WAS THIS USEFUL? 👍 👎

Malanta

Feedback ⓘ announced · Sep 16, 2026 BLOG ↗ RANK

> Malanta launched an Attribution connector for OpenCTI that turns flat `apt:` indicator labels into pivotable Intrusion Set relationships, alongside a backfill script and a native STIX/TAXII feed integration.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 STIX/TAXII 2.1 feed via built-in OpenCTI ingester NEW

90

Malanta's feed can be configured through OpenCTI's built-in ingester as STIX 2.1 / TAXII 2.1 at server URL `https://app.malanta.ai/feeds/taxii2/default`, collection `pre-attack-indicators`, using Bearer token (Malanta API key) authentication — no custom bulk-import connector required.

Wire the Malanta pre-attack feed into OpenCTI's built-in TAXII ingester so indicators arrive as STIX 2.1 objects without a custom import connector.

📍 Name: Malanta Pre-Attack Indicators TAXII server URL: `https://app.malanta.ai/feeds/taxii2/default` TAXII version: TAXII 2.1 Collection: pre-attack-indicators Authentication: Bearer <malanta-api-key> Import from date: (leave empty)

— Exact URL, collection and auth values given as a runnable setup

launch-20260916-347ce649

02 Malanta Attribution stream connector for OpenCTI NEW

73

A new Malanta Attribution stream connector, listed in the Filigran OpenCTI Integrations Library, listens to OpenCTI's live event stream and creates Intrusion Set entities plus `indicates` relationships for every Malanta indicator carrying an `apt:` label (e.g. `apt:APT44`, `apt:APT38`). Infrastructure cluster entities are preserved exactly as ingested rather than rewritten into single-actor Intrusion Set objects, keeping multi-actor staging infrastructure accurate.

— Names entity types and label scheme but no install command given

launch-20260916-347ce649

03 **Backfill script for retroactive attribution** NEW 61

A `backfill.py` script retroactively attributes indicators ingested before the connector started, running dry-run then apply, with deterministic IDs that make re-runs safe no-ops.

— Script and workflow named but no exact CLI syntax shown `launch-20260916-347ce649`

announced · Sep 16, 2026 BLOG ↗

Malanta Attribution connector for OpenCTI converts flat `apt:` labels into pivotable Intrusion Set relationships via a live event stream.

└──▷ **HOW TO FIND IT**

Wire the Malanta pre-attack feed into OpenCTI's built-in TAXII ingester so indicators arrive as STIX 2.1 objects without a custom import connector.

📍 Name: Malanta Pre-Attack Indicators TAXII server URL: `https://app.malanta.ai/feeds/taxii2/default` TAXII version: TAXII 2.1 Collection: pre-attack-indicators Authentication: Bearer <malanta-api-key> Import from date: (leave empty)

Adds a new Malanta Attribution stream connector, listed in the Filigran OpenCTI Integrations Library, that listens to OpenCTI's live event stream and creates `Intrusion Set` entities plus `indicates` relationships for every Malanta indicator carrying an `apt:` label (e.g. `apt:APT44`, `apt:APT38`).

Supports a `backfill.py` script (dry-run then apply) to retroactively attribute indicators ingested before the connector started, with deterministic IDs making re-runs safe no-ops.

Configures Malanta's STIX 2.1 / TAXII 2.1 feed via OpenCTI's built-in ingester at TAXII server URL `https://app.malanta.ai/feeds/taxii2/default`, collection `pre-attack-indicators`, using Bearer token (Malanta API key) authentication — no custom bulk-import connector required.

Preserves infrastructure cluster entities exactly as ingested — clusters are not rewritten into single-actor `Intrusion Set` objects — keeping multi-actor staging infrastructure accurate.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ announced · Sep 17, 2026 BLOG ↗ RANK

Manage devices and everything they touch. See how security, IT, and compliance get easier when unified around devices with Iru.

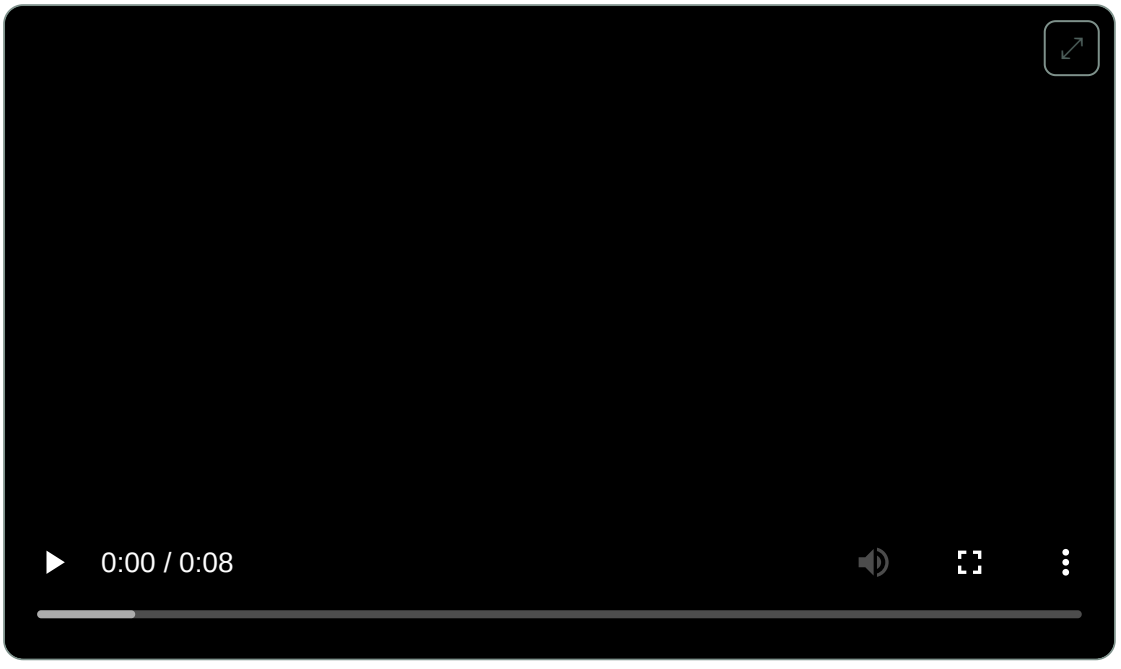
> Iru shipped a full macOS LAPS Library Item for automated local admin password rotation and escrow, plus AI-driven Windows Update configuration recommendations in the Iru chat assistant.

└──▷ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 macOS LAPS Library Item NEW

95

Adds a new macOS LAPS Library Item that automatically generates, rotates, and securely escrows local administrator passwords for Mac, manages multiple local admin accounts per device, and can import accounts from ADE Library Items and password policies from Passcode Library Items. Rotation runs on a scheduled cadence and via rotation-after-access — retiring a retrieved credential after a defined window without extending it on repeated views — through the Iru Agent using native Open Directory APIs, preserving SecureToken and Volume Ownership on Apple silicon rather than relying on Apple's `SetAutoAdminPassword` MDM command. Retrieval is gated by role-based access control, and all configuration changes, rotation events, and retrieval events (which user accessed a password, for which device, when) are captured in audit history; deployment is configured in the Library as a Library Item of type macOS LAPS and assigned via an Assignment Map.



Deploy macOS LAPS to managed Macs so every device gets a unique, automatically rotated local admin credential governed by a single central policy.

📍 In the Iru console, go to Library, create a new Library Item of type macOS LAPS, choose the local administrator account(s) to manage, configure password length and character complexity (optionally importing from a Passcode Library Item), set the rotation cadence and the rotation-after-access window, then assign the Library Item to the appropriate Assignment Map(s).

— Names exact mechanism, config surfaces, and deployment steps. launch-20260917-413408ed

Introducing macOS LAPS

02 **Windows Update recommendations in Iru AI** NEW 60

Iru AI can recommend and implement Windows Update Library Item settings based on CIS benchmarks or Microsoft defaults, asking follow-up questions to refine or explain settings, and can answer questions about Library Items, settings, and best practices from any chat without requiring navigation to a specific Library Item page.



— Describes behavior and scope but no exact command or endpoint.

Windows Update Library Item Recommendat…

announced · Sep 17, 2026 BLOG ↗

Iru: Supports automatic rotation on a scheduled cadence and rotation-after-access, so a retrieved credential is retired after a defined window without extending the window on repeated views.

└──▷ HOW TO FIND IT

Deploy macOS LAPS to managed Macs so every device gets a unique, automatically rotated local admin credential governed by a single central policy.

📍 In the Iru console, go to Library, create a new Library Item of type macOS LAPS, choose the local administrator account(s) to manage, configure password length and

character complexity (optionally importing from a Passcode Library Item), set the rotation cadence and the rotation-after-access window, then assign the Library Item to the appropriate Assignment Map(s).

Enables management and rotation of multiple local administrator accounts per device, with the ability to import accounts from ADE Library Items and password policies from Passcode Library Items.

Records detailed audit history of configuration changes, rotation events, and retrieval events — capturing which user accessed a password, for which device, and when.

Implements password rotation through the Iru Agent using native Open Directory APIs, preserving SecureToken and Volume Ownership on Apple silicon rather than relying on Apple's `SetAutoAdminPassword` MDM command.

Enforces role-based access control on password retrieval so only authorized users can view a device's local admin credential.

Introducing macOS LAPS NOTES ↗ STABLE

Iru adds macOS LAPS: automatic generation, rotation, and escrow of local administrator passwords for Mac.

Adds a new LAPS Library Item for Mac that automatically generates, rotates, securely escrows, and controls access to local administrator passwords.

Supports strong, customizable passwords with automatic rotation on a schedule and after each viewing.

Windows Update Library Item Recommendations in Iru AI NOTES ↗ LATEST STABLE

Iru AI can now recommend and apply Windows Update settings based on CIS benchmarks or Microsoft defaults via chat.

Enables Iru AI to recommend and implement Windows Update Library Item settings based on CIS benchmarks or Microsoft defaults, with follow-up questions to refine or explain settings.

Allows Iru AI to answer questions about Library Items, settings, and best practices from any chat, without requiring navigation to a specific Library Item page.

WAS THIS USEFUL? 👍 👎

Automox

Feedback

announced · Sep 16, 2026

BLOG ↗

Check out our full list of release notes, and stay up to date with what

Automox adds support for macOS 27 (Golden Gate) as a managed operating system.

WAS THIS USEFUL? 👍 👎

OSSEC HIDS

Feedback ⓘ announced · Aug 25, 2026 BLOG ↗ RANK

OSSEC is an Open Source Host-based Intrusion Detection System that performs log analysis, file integrity checking, policy monitoring, rootkit detection, real-time alerting and active response.

> OSSEC 4.3.0's headline additions are a per-agent FIM maintenance mode, TCP/TLS syslog forwarding, and ModSecurity audit log collection, backed by a wide set of new decoders/rules for other log sources and smaller agent-management and active-response improvements.

└──▷ **WHAT SHIPPED · 8 FEATURES** most completely described first ? what's the number?

01 **Per-agent FIM maintenance mode** NEW 95

Adds per-agent FIM maintenance mode via `agent_control -M` with subcommands `enable`, `disable`, `status`, and `end`, plus `-u <id>` to target a specific agent, letting file changes be accepted into the integrity baseline without generating alerts during planned maintenance windows.

Pause FIM alerting on a single agent during a planned maintenance window so file changes are silently absorbed into the baseline instead of generating floods of integrity alerts.

```
$ agent_control -M enable -u 042
```

— Exact command, subcommands and flag given with runnable example.

launch-20260825-904e5656

02 **TCP and TLS syslog output** NEW 90

`ossec-csyslogd` now supports TCP and optional TLS transport for forwarding alerts, configured via new `protocol`, `tls`, `tls_verify`, and `tls_ca` keys inside `syslog_output` in `ossec.conf`, replacing the default unencrypted UDP transport when enabled.

Forward OSSEC alerts over an encrypted TLS syslog channel to a remote SIEM, replacing the default unencrypted UDP transport.

```xml
<syslog_output>
  <server>siem.example.com</server>
  <port>6514</port>
  <protocol>tcp</protocol>
  <tls>yes</tls>
  <tls_verify>yes</tls_verify>
  <tls_ca>/var/ossec/etc/ca-bundle.crt</tls_ca>
  <format>json</format>
</syslog_output>
```

— Config keys and full XML example provided. launch-20260825-904e5656

03 **ModSecurity audit log collection** NEW 88

Adds a `modsec-audit` localfile log format to collect libmodsecurity serial audit logs, plus new decoders and rules for nginx ModSecurity error-log events, so WAF block decisions appear alongside host-level alerts.

Ingest libmodsecurity audit events into OSSEC so WAF block decisions appear alongside host-level alerts.

```xml
<localfile>
  <log_format>modsec-audit</log_format>
  <location>/var/log/modsec_audit.log</location>
</localfile>
```

— Named format with working config example included. launch-20260825-904e5656

04 **New decoders and rules for cPanel, Postfix, Snort, BIND, Windows and Asterisk**

NEW 70

Adds cPanel/cpsrvd decoders and rules for monitoring cPanel services; Postfix SMTP connect-abort detection with escalation by source IP; matching of Snort `[Drop]` fast alerts in the snort decoder; decoding of named (BIND) client IPs with an optional `@cookie` prefix; a Windows interactive logon (logon type 2) success detail rule `18262`; and an Asterisk IPv6 denied decoder so `srcip` omits square brackets.

— Six log-source additions named but no config or usage steps. launch-20260825-904e5656

05 **Agent management: bulk load and safer removal** IMPROVED 67

Supports `manage_agents -f -` to bulk-load agents from stdin, and now backs up agent state before a remove or force-delete operation while alerting on duplicated IP addresses.

— Command named; backup/alert mechanism only briefly described. launch-20260825-904e5656

THINNER COVERAGE BELOW

06 **Larger syslog alert payload buffer** IMPROVED 50

Raises the `ossec-csyslogd` alert payload buffer from 2048 bytes to `OS_MAXSTR` for default, CEF, JSON, and Splunk alert formats, avoiding truncation of larger alerts.

— Concrete numbers but nothing for the reader to configure. launch-20260825-904e5656

07 **Dedicated iptables chain for firewall-drop** IMPROVED 40

Adds a dedicated OSSEC iptables chain used by the `firewall-drop` active response, isolating its rules from other iptables chains.

— Bare description with no configuration example. launch-20260825-904e5656

08 **Trailing comments in CDB lists** IMPROVED 40

Supports `###` trailing comments in CDB list text files.

— Single-line addition with exact syntax but no context. launch-20260825-904e5656

announced · Aug 25, 2026 BLOG ↗

OSSEC 4.3.0 adds FIM maintenance mode, TCP/TLS syslog output, and ModSecurity audit log collection

└──▷ **TRY IT**

Pause FIM alerting on a single agent during a planned maintenance window so file changes are silently absorbed into the baseline instead of generating floods of integrity alerts.

```
$ agent_control -M enable -u 042
```

Forward OSSEC alerts over an encrypted TLS syslog channel to a remote SIEM, replacing the default unencrypted UDP transport.

/var/ossec/etc/ossec.conf

```
<syslog_output>
  <server>siem.example.com</server>
  <port>6514</port>
  <protocol>tcp</protocol>
  <tls>yes</tls>
  <tls_verify>yes</tls_verify>
  <tls_ca>/var/ossec/etc/ca-bundle.crt</tls_ca>
  <format>json</format>
</syslog_output>
```

Ingest libmodsecurity audit events into OSSEC so WAF block decisions appear alongside host-level alerts.

/var/ossec/etc/ossec.conf

```
<localfile>
  <log_format>modsec-audit</log_format>
  <location>/var/log/modsec_audit.log</location>
</localfile>
```

Adds per-agent FIM maintenance mode via `agent_control -M` with subcommands `enable`, `disable`, `status`, and `end` plus `-u <id>` to suppress integrity alerts while accepting file changes into the baseline.

Adds TCP and optional TLS support to `ossec-csyslogd` syslog output via new `protocol`, `tls`, `tls_verify`, and `tls_ca` config keys in `syslog_output`.

Raises the `ossec-csyslogd` alert payload buffer from 2048 bytes to `OS_MAXSTR` for default, CEF, JSON, and Splunk alert formats.

Adds `modsec-audit` localfile format to collect libmodsecurity serial audit logs, plus new decoders and rules for nginx ModSecurity error-log events.

Supports `manage_agents -f -` to bulk-load agents from stdin.

+9 more

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Dashboard enhancements and bug fixes (2026-09-17) BLOG ↗ RANK

Databricks is a unified analytics platform for data engineering, data science, and machine learning built on Apache Spark.

> Databricks' biggest updates this window are Unity Catalog governance additions — Beta metastore-level ABAC policies and automatic sensitive-data classification — alongside capacity increases for Genie Agents/Genie One and a batch of AI/BI Dashboard authoring and metric-view enhancements.

└──▷ **WHAT SHIPPED · 11 FEATURES** most completely described first ? what's the number?

01 **Live metric references in dashboard widget titles** NEW 80

Typing `@` in a visualization widget's title or description field opens a field reference picker that inserts a live, formatted field reference (e.g. sum, average) whose field, aggregation function, and number or date formatting can be changed after insertion.

Surface a live KPI directly in a widget title so viewers always see the current filtered total without reading the chart itself.

> 📍 In a visualization widget's title or description field, type @ to open the field reference picker, select the desired field and aggregation function (e.g. Sum), then confirm to insert the live metric reference.

— Names exact trigger and controls, with a runnable step example.

Dashboard enhancements and bug fixes (2…

02 **Per-facet axis scaling for faceted charts** NEW 70

Adds a Per facet axis range option on faceted charts so each subchart scales independently to its own data instead of sharing one range across all facets.

Prevent misleading axis comparisons on a faceted chart where subcharts have very different value ranges.

> 📍 Open the faceted chart's axis settings, locate Axis range, and set it to Per facet so each subchart scales to its own data.

— Exact setting named with a clear usage path. Dashboard enhancements and bug fixes (2…

03 **Metastore-level ABAC policies (Beta)** NEW 70

Unity Catalog now supports attaching ABAC policies at the metastore level (Beta), so a single row filter, column mask, GRANT, or DENY policy applies across every catalog in the metastore, including catalogs created later.

— Names policy types and scope, but no config example.

Metastore-level ABAC policies are in Be…

04 **Higher usage limits for Genie Agents** IMPROVED 60

Raises the PDF upload limit for Genie One and Genie Agents from 20 pages to 100 pages, and extends the Monitor tab so authors can summarize up to six months of Genie Agent activity, up from an unspecified shorter window.

— Concrete before/after numbers for both limits. Genie Agents enhancements and bug fixes…

05 **Multi-file and folder skills in Genie One** IMPROVED 60

Skills in Genie One now support multiple files and sub-folders, up to three folders deep and 50 files or folders per skill, limited to text files only.

— Specific numeric limits given, no usage steps. Skills in Genie One support multiple fi…

06 **Window measures on numeric index columns in metric views** NEW 60

Metric views now support window measures using a consecutive integer order column, adding unitless offset and trailing/leading range options that enable non-calendar period analysis such as fiscal weeks or 4-4-5 periods.

— Explains mechanism and use case, no example query.

Window measures on a numeric index colu…

THINNER COVERAGE BELOW

07 **Genie Agents restricted to attached data sources** BREAKING 50

A Genie Agent is now restricted to only the data assets explicitly attached as sources, preventing it from drawing on unattached data objects; any previously implicit or ambient data access will no longer work.

— Clear breaking behavior change but no migration steps.

Genie Agents enhancements and bug fixes…

08 **Genie Code support for dashboard datasets and parameters** NEW 45

Genie Code can now convert a SQL dataset into a local metric view directly from the dataset's context menu, handling the conversion automatically, and Genie Code for

dashboard authoring now understands SQL dataset parameters when building and editing dashboards.

— Describes two Genie Code additions but no UI steps.

Dashboard enhancements and bug fixes (2…

09 **Pivot table styling and layout controls** IMPROVED 45

Adds a configurable default expand level for pivot table row and column hierarchy when the table first loads, plus independent styling of pivot table totals headers and values, including conditional styling support.

— Two named pivot table controls, no walkthrough given.

Dashboard enhancements and bug fixes (2…

10 **Theme picker color palette previews** IMPROVED 23

The preset theme picker now shows color palette previews before a theme is selected.

— Single-line UI tweak with no further detail. Dashboard enhancements and bug fixes (2…

11 **Field variable option reordering** IMPROVED 18

Options within a field variable can now be reordered.

— Bare one-line description, no mechanism. Dashboard enhancements and bug fixes (2…

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! A Genie Agent now uses only data assets explicitly attached as sources; any previously implicit or ambient data access will no longer work.

Dashboard enhancements and bug fixes (2026-09-17) BLOG ↗ STABLE

Databricks AI/BI Dashboards gains metric references in widget titles, pivot-table styling, per-facet axis scaling, and Genie Code SQL parameter support.

└──▷ HOW TO FIND IT

Surface a live KPI directly in a widget title so viewers always see the current filtered total without reading the chart itself.

📍 In a visualization widget's title or description field, type @ to open the field reference picker, select the desired field and aggregation function (e.g. Sum), then confirm to insert the live metric reference.

Prevent misleading axis comparisons on a faceted chart where subcharts have very different value ranges.

📍 Open the faceted chart's axis settings, locate `Axis range`, and set it to `Per facet` so each subchart scales to its own data.

Adds the ability to convert a SQL dataset into a local metric view directly from the dataset's context menu, with Genie Code handling the conversion.

Genie Code for dashboard authoring now understands SQL dataset parameters when building and editing dashboards.

Enables inserting live, formatted field references (sum, average, etc.) into visualization widget titles and descriptions by typing `@`, with support for changing the field, aggregation function, and number or date formatting.

Adds a configurable default expand level for pivot table row and column hierarchy when the table first loads.

Adds `Per facet` axis range scaling on faceted charts, allowing each subchart to scale independently instead of sharing one range.

+3 more

Genie Agents enhancements and bug fixes (2026-09-17)

BLOG ↗ STABLE

Genie Agents gains 100-page PDF uploads, six-month activity summaries, and source-scoped data access.

Raises the PDF upload limit for Genie One and Genie Agents from 20 pages to 100 pages.

Extends the Monitor tab so authors can summarize up to six months of Genie Agent activity, up from an unspecified shorter window.

Restricts a Genie Agent to only the data assets explicitly attached as sources, preventing it from drawing on unattached data objects.

└─▷ **BREAKING ON UPGRADE**

! A Genie Agent now uses only data assets explicitly attached as sources; any previously implicit or ambient data access will no longer work.

Metastore-level ABAC policies are in Beta (2026-09-17)

BLOG ↗ STABLE

Unity Catalog gains metastore-level ABAC policies in Beta, letting one row filter, column mask, GRANT, or DENY rule span every catalog.

Skills in Genie One support multiple files and folders (2026-09-17)

BLOG ↗

LATEST STABLE

Genie One skills now support multiple files and sub-folders, up to 3 levels deep and 50 files per skill.

Window measures on a numeric index column in metric views (2026-09-14) BLOG ↗

STABLE

Databricks metric views gain window measures on numeric index columns for non-calendar periods like fiscal weeks and 4-4-5.

WAS THIS USEFUL? 👍 👎

◆ SOAR

TheHive

Feedback ⓘ announced · Sep 16, 2026 BLOG ↗ RANK

TheHive is a Collaborative Case Management Platform, now distributed as a commercial version

> TheHive 5.8 introduces native SLA tracking for alerts, cases and tasks, alongside dashboard and search improvements.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 Native SLA tracking for alerts, cases and tasks NEW

65

TheHive 5.8 adds native SLA tracking with per-entity-type rules, configurable warning thresholds, and breach notifications delivered via the existing notification system. It also supports per-organization SLA rules, letting MSSPs assign different response-time targets to each client organization.

– Explains rules, thresholds, notifications and MSSP scoping but no config keys

launch-20260916-4a5364d3

THINNER COVERAGE BELOW

02 Custom fields and tags in Dashboard Table widget IMPROVED

40

The Table widget in Dashboards is extended to display custom fields and tags alongside built-in fields.

– Names the widget but gives no configuration detail launch-20260916-4a5364d3

03 Global Search now covers comments and knowledge base IMPROVED

40

Global Search, introduced in 5.6, is expanded to cover comments and knowledge base pages in addition to existing entity types.

– States scope expansion but no mechanism or UI path launch-20260916-4a5364d3

└──▷ **ALSO FROM THESE RELEASES**

announced · Sep 16, 2026 BLOG ↗

> TheHive 5.8 adds native SLA tracking for alerts, cases and tasks, custom fields in dashboards, and expanded global search.

> Adds native SLA tracking for alerts, cases, and tasks, with per-entity-type rules, configurable warning thresholds, and breach notifications via the existing notification system.
>
> Supports per-organization SLA rules, letting MSSPs assign different response-time targets to each client organization.
>
> Extends the Table widget in Dashboards to display custom fields and tags alongside built-in fields.
>
> Expands Global Search (introduced in 5.6) to cover comments and knowledge base pages in addition to existing entity types.

WAS THIS USEFUL? 👍 👎

◆ **Detection Engineering**

sysmon-modular

Feedback ⓘ changes since 2026-08-19 CODE ↗

A repository of sysmon configuration modules

sysmon-modular tooling gains MDE config loading, path normalization, deduplication, AI/agent coverage, and versioning

└──▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-08-19
https://github.com/olafhartong/sysmon-modular.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-19
```

Adds AI and agent coverage to the detection rule set.

Adds versioning to the sysmon-modular tooling.

Adds path normalization, config areas, and deduplication to tooling processing.

Adds high-level duplicate tagging to MDE processing output.

Moves configs to releases for reproducible reporting and distribution.

+3 more

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

Qualys

Feedback ⓘ changelog · Sep 17, 2026 BLOG ↗ RANK

Cybersecurity blog covering software, services, and risk insights. Stay updated with the latest on threat detection, compliance, and vulnerability management.

> Qualys expanded Policy Audit with AI-assisted policy creation, continuous compliance monitoring, automated remediation, and risk-based prioritization of findings.

└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

01 **AI-Powered Policy Creation in Policy Audit** NEW 50

Teams can upload frameworks, benchmarks, vendor guidance, or internal policies and use AI assistance to map controls to policy logic, assessment criteria, and expected values, with human review and approval required before publication.

– Describes workflow and human-review step but no exact UI path or command.

2026-09-17-5c9e514a

02 **Audit Fix automated remediation** NEW 40

Audit Fix streamlines correction of compliance findings through pre-built remediation content and automated workflows, and validates that fixes remain in place over time.

– Names mechanism (pre-built content, validation) but no specific commands.

2026-09-17-5c9e514a

03 **Risk-based prioritization of audit findings** IMPROVED 40

Policy Audit findings are now prioritized using asset criticality, threat exposure, business context, ransomware relevance, and security risk indicators.

– Lists factors but no scoring mechanics or thresholds given. 2026-09-17-5c9e514a

04 **Audit Insights continuous compliance monitoring** NEW 35

Audit Insights continuously monitors controls, detects compliance drift, and automatically collects evidence throughout the year instead of relying on point-in-time assessments.

– Explains capability but no configuration or access detail. 2026-09-17-5c9e514a

└──▷ ALSO FROM THESE RELEASES

changelog · Sep 17, 2026 BLOG ↗ LATEST STABLE

Qualys adds AI-Powered Policy Creation, Audit Insights continuous monitoring, and Audit Fix automated remediation to Policy Audit

Adds AI-Powered Policy Creation for Policy Audit, enabling teams to upload frameworks, benchmarks, vendor guidance, or internal policies and use AI assistance to map controls to policy logic, assessment criteria, and expected values — with human review and approval before publication.

Introduces Audit Insights to continuously monitor controls, detect compliance drift, and automatically collect evidence throughout the year instead of at point-in-time assessments.

Introduces Audit Fix to streamline correction of compliance findings through pre-built remediation content and automated workflows, validating that fixes remain in place over time.

Adds risk-based prioritization to Policy Audit findings by incorporating asset criticality, threat exposure, business context, ransomware relevance, and security risk indicators.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Axonius

Feedback ⓘ Axonius Release Notes 9.0.7 (2026-09-15) BLOG ↗

RANK

Axonius is a cybersecurity platform that discovers, manages, and protects all devices and software across an organization's IT infrastructure.

> Axonius 9.0.7 introduces a CMDB Reconciliation workspace for automated, audit-ready CMDB syncing, adds Windows patch-release enrichment fields and 11 new adapters, and rolls out a Network Inspectors management page, a Custom Data Asset Source mechanism, and numerous adapter-level enhancements, alongside a breaking change to how Citrix ADC NAT rules are modeled as assets.

└──▷ **WHAT SHIPPED · 18 FEATURES** most completely described first ? what's the number?

01 **Windows patch release enrichment fields** NEW 95

Adds `OS: Windows Patch Release Type` and `OS: Windows Patch Release Date` enrichment fields, nested under the OS complex field on Windows device records, updated daily from the Microsoft Windows Update catalog, enabling queries such as assets where patch release date is more than X days ago.

Find Windows devices where the patch is more than 30 days old, using the new OS: Windows Patch Release Date enrichment field, to prioritize remediation of lagging endpoints.

📍 In the Devices page, create a new query with the filter: OS: Windows Patch Release Date is more than 30 days ago, then save and export the results to drive your patch compliance report.

— Exact field names, update cadence, and a runnable query example

Axonius Release Notes 9.0.7 (2026-09-15)

02 **Druva Cloud Platform adapter enhancements** NEW 95

Adds a Devices API Version connection setting to the Druva Cloud Platform adapter, allowing selection between version 1 (default) and version 2 of the Druva Devices API, and adds Application Settings support fetching administrator roles and statuses from the Druva administrators API and evaluating audit event availability.

Switch the Druva adapter to the v2 Devices API to retrieve devices using the updated endpoint.

📍 In the Druva Cloud Platform adapter connection settings, set Devices API Version to version 2, then enable admin APIs in the Druva Console under Settings ›

Enable admin APIs, and trigger a fetch to populate updated device and Application Settings assets.

— Exact setting name, versions, and full switch-over steps given

Axonius Release Notes 9.0.7 (2026-09-15)

03 **Confluent adapter Application Settings collection** NEW

80

Adds Fetch Application Settings advanced setting (default: disabled) to the Confluent adapter, which fetches Application Settings assets from Confluent cluster ACLs when enabled.

Enable collection of Application Settings from a Confluent cluster to gain visibility into ACL-based access control posture.

📍 In the Confluent adapter configuration, navigate to Advanced Settings and enable Fetch Application Settings (default: disabled), then save and run a fetch cycle to populate Application Settings assets.

— Exact flag name, default, and enable steps provided

Axonius Release Notes 9.0.7 (2026-09-15)

04 **Citrix ADC NAT assets now modeled as Load Balancers** BREAKING

70

The Citrix ADC adapter no longer publishes Network Route assets for NAT rules; NAT translations now appear on Load Balancer assets instead. Existing queries, reports, or enforcement workflows targeting Network Route assets from Citrix ADC NAT must be updated to use NAT translations on Load Balancer assets.

— Clear before/after and migration guidance, no exact query syntax

Axonius Release Notes 9.0.7 (2026-09-15)

05 **11 new adapters added** NEW

65

Adds 11 new adapters: DTS Identity (Users, Devices, Aggregated Security Findings, SaaS Applications, Software), HPE Alletra SAN (Devices), Moody's CreditLens (Users), NetSPI Resolve (Devices), OneSpan (Users, Groups, Application Settings), Radware Cloud Services (Networks, Application Settings), SOCRadar (Devices, Networks, URLs, Certificates), Spacelift (Users, Application Settings), Ubiquiti UISP (Devices), UiPath (Users, Application Settings), and WordPress Toolkit (URLs, Aggregated Security Findings, SaaS Applications).

— Every adapter and asset type named but no setup detail

Axonius Release Notes 9.0.7 (2026-09-15)

06 **DocuSign multi-account collection** NEW 65

Adds Single Account / All Organization Accounts selection to the DocuSign adapter, with a new Organization ID field to discover and collect data across all accounts in a DocuSign organization.

— Exact fields named but no step-by-step configuration example

Axonius Release Notes 9.0.7 (2026-09-15)

07 **Network Inspectors management page** NEW 60

Adds a new Network Inspectors page under Settings → System for Cyber-Physical Assets deployments, providing centralized management of Virtual Network Inspector appliances including add, token generation, token re-generation, and removal.

— Exact navigation path given but no step-by-step example

Axonius Release Notes 9.0.7 (2026-09-15)

08 **Custom Data Asset Source mechanism** NEW 60

Adds a Custom Data Asset Source mechanism with a dedicated enforcement action that converts manually-entered custom data into a first-class adapter contribution, processed by the correlation engine with no defined retention period.

— Mechanism described but no exact configuration steps

Axonius Release Notes 9.0.7 (2026-09-15)

09 **AWS adapter CloudFront alias discovery** NEW 60

Adds support for fetching Domains & URLs for CloudFront distribution aliases in the AWS adapter, linking each alias to its CloudFront distribution for inclusion in Internet Exposure queries.

— Names the use case (Internet Exposure queries) but no exact steps

Axonius Release Notes 9.0.7 (2026-09-15)

10 **AppViewX CERT adapter Application Settings** NEW 60

Adds Application Settings asset type support to the AppViewX CERT adapter, fetching authentication, access control, credential, key rotation, logging, retention, login, SMTP, SSH connection, and service account settings.

— Lists all settings categories but no navigation or config steps

Axonius Release Notes 9.0.7 (2026-09-15)

THINNER COVERAGE BELOW

11 **Delinea Vault Is File setting** NEW 55

Adds an optional Is File setting to Delinea Privilege Manager (Thycotic) Vault secret references, causing Axonius to retrieve file content instead of the field value when enabled.

– Exact setting name and behavior, no full workflow example

Axonius Release Notes 9.0.7 (2026-09-15)

12 **Datadog Lambda host reclassification** IMPROVED 55

Reclassifies Datadog host records tagged with `aws:lambda` or `functionname:` as Serverless Function assets instead of Device assets.

– Exact tags named but change is automatic, no user action

Axonius Release Notes 9.0.7 (2026-09-15)

13 **CMDB Reconciliation workspace** NEW 50

A new CMDB Reconciliation workspace automatically aggregates, validates, and pushes data to CMDBs with zero manual data entry, supporting audit-readiness and licensing compliance.

– Names the workspace but no UI path or setup steps given

Axonius Release Notes 9.0.7 (2026-09-15)

14 **Dynatrace adapter fortinet:fortigate mapping** NEW 45

Adds support for the `fortinet:fortigate` entity type in the Dynatrace adapter, mapping those entities to Device assets in Axonius.

– Entity type named but no configuration or usage detail

Axonius Release Notes 9.0.7 (2026-09-15)

15 **Optional endpoint fields in Remediation Ownership** IMPROVED 45

Makes destination endpoint fields in Remediation Ownership configuration optional, allowing remediation owners to be assigned to Security Findings without providing endpoint details.

– Describes the change but no exact field names or steps

Axonius Release Notes 9.0.7 (2026-09-15)

16 **Earliest First Fetch Time field for vulnerabilities** NEW 40

Adds a new Earliest First Fetch Time field to the Vulnerability Repository and Aggregated Security Findings pages, recording the earliest first-fetch timestamp across all sources for each vulnerability.

– Field and pages named but no usage steps Axonius Release Notes 9.0.7 (2026-09-15)

17 **Cynerio upper-bound detection date setting** NEW 40

Adds an upper-bound detection date advanced setting to the Cynerio adapter for more granular control over the time range when fetching risks and vulnerabilities.

– No exact field name or steps to configure the setting

Axonius Release Notes 9.0.7 (2026-09-15)

18 **Default view per role for Assets pages** NEW 40

Enables admins to set a default view per role for specific Assets pages, automatically applied when a user with that role opens the page.

– No specific page names or configuration path given

Axonius Release Notes 9.0.7 (2026-09-15)

└──▷ **ALSO FROM THESE RELEASES**

└──▷ **BREAKING ON UPGRADE**

! The Citrix ADC adapter no longer publishes Network Route assets for NAT rules; NAT translations now appear on Load Balancer assets instead. Existing queries, reports, or enforcement workflows targeting Network Route assets from Citrix ADC NAT must be updated to use NAT translations on Load Balancer assets.

Axonius Release Notes 9.0.7 (2026-09-15) BLOG ↗ LATEST STABLE

Axonius 9.0.7 adds CMDB Reconciliation workspace, 11 new adapters, Signature Mitigations for CVE findings, and Virtual Network Inspector management.

└──▷ **HOW TO FIND IT**

Find Windows devices where the patch is more than 30 days old, using the new OS: Windows Patch Release Date enrichment field, to prioritize remediation of lagging endpoints.

📍 In the Devices page, create a new query with the filter: OS: Windows Patch Release Date is more than 30 days ago, then save and export the results to drive your patch compliance report.

Enable collection of Application Settings from a Confluent cluster to gain visibility into ACL-based access control posture.

📍 In the Confluent adapter configuration, navigate to Advanced Settings and enable Fetch Application Settings (default: disabled), then save and run a fetch cycle to populate Application Settings assets.

Switch the Druva adapter to the v2 Devices API to retrieve devices using the updated endpoint.

📍 In the Druva Cloud Platform adapter connection settings, set Devices API Version to version 2, then enable admin APIs in the Druva Console under Settings › Enable

admin APIs, and trigger a fetch to populate updated device and Application Settings assets.

Adds a new `CMDB Reconciliation` workspace that automatically aggregates, validates, and pushes data to CMDBs with zero manual data entry, supporting audit-readiness and licensing compliance.

Adds `OS: Windows Patch Release Type` and `OS: Windows Patch Release Date` enrichment fields (nested under the OS complex field) to Windows device records, updated daily from the Microsoft Windows Update catalog, enabling queries such as 'assets where patch release date is more than X days ago'.

Adds a new `Earliest First Fetch Time` field to the Vulnerability Repository and Aggregated Security Findings pages, recording the earliest first-fetch timestamp across all sources for each vulnerability.

Adds a new `Network Inspectors` page under `Settings → System` for Cyber-Physical Assets deployments, providing centralized management of Virtual Network Inspector appliances including add, token generation, token re-generation, and removal.

Adds a `Custom Data Asset Source` mechanism with a dedicated enforcement action that converts manually-entered custom data into a first-class adapter contribution, processed by the correlation engine with no defined retention period.

`+13 more`

└─▷ **BREAKING ON UPGRADE**

! The Citrix ADC adapter no longer publishes Network Route assets for NAT rules; NAT translations now appear on Load Balancer assets instead. Existing queries, reports, or enforcement workflows targeting Network Route assets from Citrix ADC NAT must be updated to use NAT translations on Load Balancer assets.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta

Feedback ⓘ snapshot-20260918 NOTES ↗ RANK

Use Immuta to provision governed access to data.

> Immuta added a new Teradata connector with full policy and audit support, expanded Redshift connection management, and let the Request app handle multi-user access requests in one submission, while announcing the phase-out of Snowflake username/password authentication and fingerprint-based masking.

└──▷ **WHAT SHIPPED · 5 FEATURES** most completely described first ? what's the number?

01 **Snowflake fingerprint-based masking removal** DEPRECATED

75

Fingerprint-based masking options for Snowflake (format-preserving masking, randomized response masking, and rounding/grouping using a data fingerprint) are deprecated and will be removed from Immuta SaaS on December 15, 2026; tenants without existing policies using these options can no longer create them.

— Names exact masking types and removal date snapshot-20260918

02 **Snowflake username/password authentication deprecated**

70

BREAKING

New Snowflake connections can no longer be created using username/password authentication, following Snowflake's deprecation of that method for service users; existing connections continue to work until Snowflake completes its deprecation phases in October 2026.

— Gives before/after and exact deprecation timeline snapshot-20260918

03 **Multi-user access requests in Request app** NEW

60

The Request app now supports requesting access for multiple users simultaneously in a single submission, replacing the need to submit separate requests per person. Hover-over tooltips on user names now display each user's email address to help distinguish users with similar or identical names.

When onboarding a team to a new dataset, submit a single governed access request covering all members instead of repeating the process per person.

📍 In the Request app, open the access request form, select all relevant users from the user picker, fill in the request details, and click Submit once to create a single request covering all selected users.

— Includes concrete workflow steps for submitting a group request snapshot-20260918

THINNER COVERAGE BELOW

04 **Teradata data source integration** NEW 55

Adds Teradata integration with policy enforcement, access control, authentication, and audit support for data in Teradata environments.

— Names capabilities covered but no setup steps given product docs

05 **Amazon Redshift connection management settings** IMPROVED 40

Adds Amazon Redshift connection management settings, including the ability to remove a Redshift connection.

— Thin description of a single added setting product docs

└──▷ BREAKING ON UPGRADE

! New Snowflake connections can no longer be created using username/password authentication, following Snowflake's deprecation of that method for service users; existing connections continue to work until Snowflake completes its deprecation phases in October 2026.

! Fingerprint-based masking options for Snowflake (format-preserving masking, randomized response masking, and rounding/grouping using a data fingerprint) are deprecated and will be removed from Immuta SaaS on December 15, 2026; tenants without existing policies using these options can no longer create them.

snapshot-20260918 NOTES ↗ LATEST STABLE

Immuta Request app now supports multi-user access requests in a single submission

└──▷ HOW TO FIND IT

When onboarding a team to a new dataset, submit a single governed access request covering all members instead of repeating the process per person.

📍 In the Request app, open the access request form, select all relevant users from the user picker, fill in the request details, and click Submit once to create a single request covering all selected users.

Supports requesting access for multiple users simultaneously in the Request app, replacing the need to submit separate requests per person.

Adds hover-over tooltips on user names in the Request app to display each user's email address, helping distinguish users with similar or identical names.

└──▷ **BREAKING ON UPGRADE**

! New Snowflake connections can no longer be created using username/password authentication, following Snowflake's deprecation of that method for service users; existing connections continue to work until Snowflake completes its deprecation phases in October 2026.

! Fingerprint-based masking options for Snowflake (format-preserving masking, randomized response masking, and rounding/grouping using a data fingerprint) are deprecated and will be removed from Immuta SaaS on December 15, 2026; tenants without existing policies using these options can no longer create them.

WAS THIS USEFUL?  

Varonis

Feedback ⓘ announced · Sep 15, 2026 BLOG ↗ RANK

The leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis launched Data Lifecycle Management, adding a unified ROT dashboard, configurable lifecycle policies with reversible quarantine, compliance-based retention enforcement, and detailed duplicate-group views.

└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

01 Retention enforcement mapped to compliance windows NEW

70

Retention enforcement maps classification categories to legal retention windows — financial and PCI data over seven years, PII and GDPR data over two years, PHI over six years, and education data over five years — and quarantines data that exceeds those windows.

— Specific compliance windows given but no configuration steps. launch-20260915-e5a43d40

02 Unified ROT dashboard with cost-savings view NEW

65

A new dashboard shows total redundant, obsolete, and trivial (ROT) data volume across all connected platforms, filterable by platform, classification category, sensitivity, data source, file type, and owner. It also surfaces storage spend and compliance risk exposure to prioritize cleanup efforts.

— Names six filter dimensions but no exact UI path or numbers. launch-20260915-e5a43d40

03 Duplicate group detail views NEW

60

New detail views for duplicate groups show which copy is the original, which is most recent, where copies live across platforms, and whether each copy is sensitive, overexposed, or past retention.

— Describes drill-in view content but not exact navigation steps.

launch-20260915-e5a43d40

THINNER COVERAGE BELOW

04 Configurable lifecycle policies with reversible quarantine

NEW

47

Out-of-the-box, fully configurable lifecycle policies for duplicates and staleness automatically move violating data to a secure, reversible quarantine.

— Mechanism named but no config keys or thresholds given. launch-20260915-e5a43d40

announced · Sep 15, 2026 BLOG ↗

Varonis: Adds a unified ROT dashboard showing total redundant, obsolete, and trivial (ROT) data volume across all connected platforms, filterable by platform, classification category, sensitivity, data source, file type, and owner.

Introduces out-of-the-box, fully configurable lifecycle policies for duplicates, staleness, and retention that automatically move violating data to a secure, reversible quarantine.

Adds retention enforcement that maps classification categories to legal windows — financial and PCI data over seven years, PII and GDPR data over two years, PHI over six years, and education data over five years — and quarantines data that exceeds those windows.

Surfaces duplicate group detail views showing which copy is the original, which is most recent, where copies live across platforms, and whether they are sensitive, overexposed, or past retention.

Enables cost-savings prioritization in the ROT dashboard, surfacing storage spend and compliance risk exposure to guide cleanup efforts.

WAS THIS USEFUL? 👍 👎

CrowdStrike Falcon Data Protection

Feedback ⓘ

undated-1f3abcdb BLOG ↗ RANK

CrowdStrike Falcon Data Protection provides endpoint detection and response capabilities to protect against threats and data breaches.

> CrowdStrike added on-device AI-powered sensitive data classification to Falcon Data Security, running locally on endpoint hardware across Windows, Linux/cloud, and macOS to catch unstructured sensitive data without cloud inference.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **On-device AI data classification across platforms** NEW

68

Falcon Data Security now runs AI-powered sensitive data classification locally rather than in the cloud: on Windows via Intel Core Ultra AI PCs using Intel's NPU, on Falcon Data Security for Cloud via Linux-based EC2 instances with NVIDIA GPUs through an ONNX Runtime Execution Provider swap, and on macOS via CoreML targeting Apple's Neural Engine (ANE) for real-time inference. This enables context-aware detection of unstructured sensitive data — such as credentials shared in natural language, chat logs, AI prompts, and free-text fields — that rule-based pattern matching cannot cover, and keeps data from leaving the endpoint.

— Names hardware/runtime specifics but no user-facing config or command undated-1f3abcdb

> undated-1f3abcdb BLOG ↗
>
> Falcon Data Security gains on-device AI classification using Intel NPU hardware, enabling context-aware sensitive data detection without cloud inference.
>
> Adds on-device AI-powered sensitive data classification to Falcon Data Security, running language models locally on Intel Core Ultra AI PCs via Intel's NPU — no data leaves the endpoint.
>
> Extends the same on-device AI classification model to Falcon Data Security for Cloud on Linux-based EC2 instances with NVIDIA GPUs via an ONNX Runtime Execution Provider swap.
>
> Brings on-device AI classification to macOS through CoreML, targeting Apple's Neural Engine (ANE) directly for real-time, on-device inference.
>
> Enables context-aware detection of unstructured sensitive data — such as credentials shared in natural language, chat logs, AI prompts, and free-text fields — that rule-based pattern matching cannot cover.

WAS THIS USEFUL?  

◆ Identity & Access Management

JumpCloud

Feedback ⓘ Release Notes 2022 (2026-09-18) BLOG ↗

RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> Across nine releases spanning 2022 through 2026, JumpCloud grew from launching Password Manager and Remote Assist into a full Zero Trust and unified IT platform, culminating in the GA of Password Vault, generally-available Access Risk Detection with behavior-based anomaly detection, Agent Identities for non-human IAM, and Device Monitoring and Alerting, alongside major expansions to RADIUS, JumpCloud Go, Android/Apple/Windows MDM, Asset Management, and the Multi-Tenant Portal.

└──▷ **WHAT SHIPPED · 49 FEATURES** most completely described first ? what's the number?

01 Pipe character escaping in systemusers API filters IMPROVED

85

Adds `|` (pipe) escape support via `\` in `in` and `nin` API filter operators, enabling pipe characters within search terms — e.g. `GET /systemusers?filter=department:$in:Finance | Accounting|IT`.

Query users whose department is `Finance | Accounting` or `IT`, using the newly escapable pipe character in an `in` filter.

```
$ curl -X GET 'https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT' \
  -H 'x-api-key: <your-api-key>'
```

— Exact operator syntax with a runnable curl example. Release Notes 2024 (2026-09-18)

02 Device Monitoring and Alerting NEW

85

Introduces JumpCloud Device Monitoring and Alerting — monitor device health (availability, software, uptime, disk, battery), policy/command/software management results, directory events (user addition, privilege elevation), and custom alerts via command execution monitoring, manageable from the Admin Portal or MTP; supports Windows, macOS, and Linux.

— Broad monitored surfaces and portals named, no step-by-step.

Release Notes 2024 (2026-09-18)

03 **JumpCloud Password Vault** NEW 85

Adds Password Vault to the Admin Portal — enable the vault, assign access to users and user groups, manage shared credentials and websites with Connect, View, Edit, and Manage permissions, monitor password health from an Overview dashboard, and view audit logs in Directory Insights — and to the User Portal, supporting credential and website management, shared folders, and bulk import via CSV or Excel. Adds Password Vault autofill to the JumpCloud Go browser extension for non-SSO website password autofill, 2FA code launch, secret viewing, and save-on-login/update/signup without device registration. At GA, adds a Password Health Dashboard covering expired, expiring, weak, and inactive credentials with a built-in password generator; four sharing permission levels — Connect, View Secret, View Details, and Manage; unification of Password Vault, Go authentication, SSO access, and SaaS management in one browser extension; JumpCloud Protect mobile app combining corporate MFA and Vault access in one app; User Credential Transfer for reassigning an offboarded employee's credentials via an orphaned folder; and immutable audit logs via Directory Insights.

Review the Password Health Dashboard to proactively identify weak or expired credentials across the organization before a breach forces a reactive response.

> 📍 In the JumpCloud Admin Portal, go to Password Vault › Password Health Dashboard to view organization-wide credential health, including expired, expiring, weak, and inactive credentials.

Reassign a departing employee's credentials to a colleague to prevent business disruption during offboarding.

> 📍 In the JumpCloud Admin Portal, navigate to Password Vault, locate the orphaned folder created for the offboarded user, and assign it to the replacement user.

— Comprehensive named surfaces across portals, extension, and mobile app.

Release Notes 2026 (2026-09-18) launch-20260915-8f4a306d

---

04 **Disk encryption device condition in Conditional Access Policies**

NEW 80

Adds a disk encryption device condition to Conditional Access Policies — admins can allow or deny access based on FileVault (macOS), system drive encryption (Windows), or root disk encryption (Linux) status. Configured via Admin Portal > Security > Conditional Access by adding a Device Condition and setting the encryption requirement per OS, with a Deny action for unencrypted devices.

Enforce a conditional access policy that blocks login from any device without full-disk encryption — useful for meeting compliance requirements across mixed OS fleets.

📍 In the Admin Portal, go to Security › Conditional Access, create or edit a policy, add a Device Condition, and set the encryption requirement to `FileVault enabled` (macOS), `system drive encrypted` (Windows), or 'root disk is encrypted' (Linux), then set the policy action to Deny.

— Exact navigation path and per-OS setting are given. Release Notes 2022 (2026-09-18)

---

05 **JumpCloud Go passwordless authentication and session management**

NEW 80

JumpCloud Go browser extension enables hardware-protected, phishing-resistant passwordless login using Apple Touch ID and Windows Hello on managed macOS and Windows devices; enforces user and device verification on every SSO session established via Go, with conditional access policies now defaulting to JumpCloud Go as the MFA method. Later adds a `User Portal > Security > Multi-Factor Authentication` section to manage JumpCloud Go sessions on user devices, and sessions now automatically refresh on device login or unlock.

— Broad mechanism and session behaviour detailed across releases.

Release Notes 2024 (2026-09-18) Release Notes 2023 (2026-09-18)

---

06 **Password Manager Desktop App auto-lock and deployment controls**

NEW 80

Introduces an auto-lock schedule with a default of 15 minutes of inactivity and a maximum of 1,440 minutes (24 hours), an Auto Lock setting that locks the app when the device is locked, and admin ability to install and update the desktop app via a new command template v.2.0.0 on Windows and macOS, removing the need for user admin permissions after the first run.

— Exact thresholds and command template version named.

What's New: Password Manager Desktop Ap…

---

07 **Multi-Tenant Portal management, commands, and billing visibility**

IMPROVED 75

MSP admins can create and reuse custom command templates across tenant organizations, and a new MSP homepage widget displays an aggregate view of all tenant reports available in the last 24 hours with bulk export to JSON or CSV. Admins can run ad hoc commands and view command-to-device bindings from both the Multi-Tenant Portal and Admin Portal, add users and devices, export user and device lists, and perform bulk actions directly from the Multi-Tenant Portal. Adds MSP billing visibility for committed-user plans, including available and consumed license counts and JumpCloud contract

download, plus a Global MFA Requirement setting under `Account > Settings` for bulk MFA enforcement across all MSP admins.

Enforce MFA for all MSP admins in bulk from the Multi-Tenant Portal to close gaps left by per-account settings.

> 📍 In the Admin Portal, go to Account › Settings and enable the Global MFA Requirement toggle under the Global MFA Requirement section.

— Many named MTP surfaces across releases, one example. Release Notes 2024 (2026-09-18) Release Notes 2023 (2026-09-18) Release Notes 2022 (2026-09-18)

---

08 **Password Manager Desktop App platform and capacity expansion**

IMPROVED 75

Adds a Password Manager Desktop App build for Linux (Debian/Ubuntu x86) with automatic updates. Later adds region selection with EU region support on the activation page (desktop app) and pairing page (browser extension); scales item storage to up to 30,000 items (without autofill) or 13,000 items (with autofill), with a single shared folder holding up to 5,000 items; adds an ARM build for macOS M1-and-above silicon with automatic upgrade from x64; and adds an RPM build for Linux (x64 Red Hat/Fedora), both auto-updating.

— Concrete limits and builds named across two releases.

What's New: Password Manager Desktop Ap… Release Notes 2023 (2026-09-18)

---

09 **Cross-platform browser and OS patch management policy and reporting**

NEW 70

Introduces a universal Google Chrome Browser Patch Management policy covering macOS, Windows, and Linux simultaneously, plus the OS Patch Management Policy Report (showing applied policies and status per device) and the Browser Patch Management Policy Report (auditing device browser versions across managed devices).

Keep Google Chrome patched across every managed OS without creating per-platform policies — ideal for orgs running mixed macOS, Windows, and Linux fleets.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy, select Universal Browser Patch Management, choose Google Chrome as the target browser, configure the patch schedule, and assign the policy to the desired device group.

— Names three surfaces and includes navigation steps. Release Notes 2022 (2026-09-18)

---

10 **Android device management (EMM)** NEW 70

Supports Android zero-touch enrollment for fully managed, dedicated, and company-owned Android devices with a work profile; enables Android EMM support for fully

managed and dedicated devices for all customers, releasing 7 new policies including Factory Reset Protection, Kiosk Mode, and System Updates; adds support for enrolling and managing company-owned and personal Android mobile devices via Android EMM work profile; and adds full Android 15 MDM support.

— Names enrollment modes and seven policies, no steps. Release Notes 2024 (2026-09-18) Release Notes 2023 (2026-09-18)

11 **Password Manager admin controls for backups and shared folders**

IMPROVED 70

Password Manager Cloud Backups let admins enable a secure cloud-based fallback to minimize data loss, and admins can manage Password Manager shared folder assignments and permissions from the Admin Portal with manual per-folder sync from the desktop app. Later adds the ability for admins to create, update, and manage shared folders — including assigning users or user groups to any folder — a Move All items button to relocate all items from a folder in one action, and shared folder events synced to Directory Insights in the Admin Portal.

— Several named admin controls added across two releases.

What's New: Password Manager Desktop Ap… Release Notes 2023 (2026-09-18)

12 **JumpCloud Protect push authentication enhancements** IMPROVED

70

Adds Push Bombing Mitigation to JumpCloud Protect, blocking more than one push notification per resource within a 60-second window, configurable in MFA Configurations; enables push notifications to be approved or denied directly from the device lock screen via long-press (iOS) or notification expansion (Android); and adds user group exclusion from JumpCloud Protect Push enforcement, allowing excluded groups to use TOTP only.

Protect admin accounts against MFA push bombing by tuning the concurrent attempt limit in MFA Configurations.

> 📍 In the Admin Portal, go to Security › MFA Configurations › JumpCloud Protect, locate Push Bombing Mitigation, and adjust the maximum concurrent attempts limit or disable the 60-second block window to match your organisation's risk tolerance.

— Named mechanisms plus a configuration example. Release Notes 2025 (2026-09-18) Release Notes 2023 (2026-09-18)

13 **Windows App Catalog** NEW

70

Introduces the JumpCloud App Catalog for Windows, enabling install and auto-update management of 1Password, Adobe Acrobat Reader DC, Chrome, Firefox, Google Drive, OneDrive, OpenVPN, Slack, Webex, and Zoom.

— Ten named apps but no setup steps given. Release Notes 2024 (2026-09-18)

14 **JumpCloud Access Requests** NEW 70

Introduces JumpCloud Access Requests, enabling configurable approval flows, in-console request management, Directory Insights event tracking, and webhook integration with ticketing systems such as JIRA.

— Named workflow and webhook integration, no setup walkthrough.

Release Notes 2025 (2026-09-18)

15 **Admin role management: Asset Manager and Custom Admin Roles**

NEW 70

Introduces a new Asset Manager role in the Admin Portal for dedicated access to Asset Management without broader Admin Portal permissions, and introduces Custom Admin Roles for granular, responsibility-based role configuration across direct organizations.

Scope a procurement team member's admin access to Asset Management only, without granting broader Admin Portal rights.

📍 In the Admin Portal, go to Administrators › Roles, assign the built-in Asset Manager role to the admin account for the procurement team member.

— Two named roles with a setup example. Release Notes 2025 (2026-09-18)

16 **MCP server discovery in AI & SaaS Management** NEW 70

Adds an `Access > AI & SaaS Management > Applications` view with an MCP Connection column showing Discovered, Connected, or Not Available status for each app, plus an Add MCP Server action to connect discovered MCP servers to AI Gateway without switching tools.

Surface and connect a device-discovered MCP server to AI Gateway from the SaaS Management view, without switching to a separate tool.

📍 1. Go to Access › AI & SaaS Management › Applications. 2. Check the MCP Connection column for an app showing Discovered. 3. Click the status to review the discovered MCP servers. 4. Click Add MCP Server to register it in AI Gateway.

— Named UI path, statuses, and a worked example. Release Notes 2026 (2026-09-18)

17 **RADIUS and network authentication enhancements** NEW

65

JumpCloud RADIUS gains certificate-based (passwordless) authentication supporting any X.509-compliant certificate authority, letting organizations import existing certificates, and authentication via Azure AD credentials so Azure AD can serve as the identity provider for RADIUS without managing users separately. Adds faster EAP

reconnect, letting recently authenticated devices skip the full security handshake using a cached session credential.

— Names mechanisms across two releases, no exact configuration steps.

Release Notes 2025 (2026-09-18) Release Notes 2022 (2026-09-18)

18 **New MDM policy types across Windows and Apple** NEW

65

Adds six new Windows MDM policies via the CSP framework: Windows Wi-Fi Configuration Policy, Windows VPN Policy, Install Certificate Policy, Bluetooth Configuration Policy, Control Policy Conflict, and Windows Notification Service, plus a policy to configure Simple Certificate Enrollment Protocol (SCEP) for Windows devices. Later adds two MDM policies restricting Apple Intelligence features on iOS/iPadOS 18.1 and macOS 15.1 Sequoia devices, and three new Windows Personalization policies: Taskbar Customization, Screen Saver Configuration, and Desktop and Lock Screen Settings.

— Twelve named policies across two releases, each thin. Release Notes 2025 (2026-09-18)

Release Notes 2024 (2026-09-18)

19 **Remote Assist launch** NEW

60

Launches JumpCloud Remote Assist for free remote desktop access on Windows and macOS devices (requires JumpCloud Agent v1.61.0 or later), with no limits on session time or number of sessions.

— Version requirement and limits stated, no walkthrough. Release Notes 2022 (2026-09-18)

20 **Remote Assist session controls and architecture-specific packages**

IMPROVED

60

Remote Assist adds access to the initial Windows Server login screen (admins can authenticate and interact as any user), user account switching during sessions, and a Ctrl+Alt+Del command shortcut in the Session Options toolbar. Later ships as two architecture-specific packages — ARM64 for Apple Silicon and x86_64 for Intel — automatically selecting the appropriate binary.

— Named session features across releases, no steps. Release Notes 2024 (2026-09-18)

Release Notes 2023 (2026-09-18)

21 **Password Manager Desktop App UI quality-of-life improvements**

IMPROVED

60

Adds Download Logs and Keyboard Shortcuts options under Settings, multi-select functionality, and item sorting; later adds archiving (view, delete, and restore archived items), pagination displaying 100 items per page, a Copy button in the Password History area, sorting items by nickname, and a What's New link on Settings > General.

— Several named UI additions with little mechanism.

What's New: Password Manager Desktop Ap… | Release Notes 2024 (2026-09-18)

22 **LDAP cipher suite restriction** BREAKING 60

LDAP no longer accepts connections using CBC-mode or SHA-1/SHA-256 integrity cipher suites with RSA or ECDHE key exchange — any LDAP client configured to use those suites will fail to connect.

— Names the exact dropped cipher/key-exchange combinations.

Release Notes 2025 (2026-09-18)

THINNER COVERAGE BELOW

23 **Password Manager launch with encrypted sync and sharing** NEW 55

Launches JumpCloud Password Manager with end-to-end encrypted local sync, browser extension auto-fill, and secure password sharing — no master password required.

— Mechanism named but no setup path. Release Notes 2022 (2026-09-18)

24 **SaaS Management app discovery** NEW 55

Introduces SaaS Management — discover and track SaaS app use via the JumpCloud Go browser extension, categorize apps, restrict access with warning or blocking messaging, and gather usage insights.

— Mechanism described but no configuration steps. Release Notes 2024 (2026-09-18)

25 **Asset Management enhancements** IMPROVED 55

Adds asset cloning with auto-clearing of unique fields to prevent duplicates, file upload support (Invoice field under Purchase Details, plus custom File Upload type fields) for attaching invoices or contracts, Disk Size added to Hardware Specifications, and Local Device Accounts added to the Asset Detail page.

— Several named fields, no walkthrough. Release Notes 2025 (2026-09-18)

26 **Directory Insights filter refinement for application and resource names**

IMPROVED 55

Refines Directory Insights filters, splitting Application Name into Application Name , Application Display Name , and Application Display Label , and Resource Name into Resource Name , Resource Display Name , and Resource Display Label .

— Named field splits, no worked example. Release Notes 2025 (2026-09-18)

27 **JumpCloud Case Portal** NEW 50

Launches the JumpCloud Case Portal, centralizing support cases and ideas from all organization admins into a filterable list accessible from the Support menu.

— Names the nav location but thin on mechanism. Release Notes 2022 (2026-09-18)

28 **Commands queue enrollment automation and running status** NEW

50

Adds a Command After Agent Install command type so admins can specify enrollment commands to run immediately after the JumpCloud agent is installed on new devices, and adds a Running column to the commands queue to distinguish commands waiting in the queue from commands actively running.

— Names two surfaces with no walkthrough. Release Notes 2022 (2026-09-18)

29 **New SSO and remote access integrations: Deel and AWS Verified Access**

NEW

50

Adds the Deel SSO Connector integration for JumpCloud SSO, and adds AWS Verified Access integration using JumpCloud Go, SAML, or OIDC SSO as a VPN-less remote access solution.

— Two named integrations, no setup detail. Release Notes 2023 (2026-09-18)

30 **Apple MDM enrollment enhancements** NEW

50

Adds Account-Driven User Enrollment (ADUE) for BYOD iOS/iPadOS 18+ devices in MDM, now in general availability, and adds Apple Business Manager as a new service discovery method for Account-Driven Enrollment in Apple MDM.

— Two named enrollment methods with thin mechanism. Release Notes 2025 (2026-09-18) Release Notes 2024 (2026-09-18)

31 **Identity First Login** NEW

45

Introduces Identity First Login for the User Portal and SSO applications — a two-screen flow that separates username/email entry from password entry.

— Describes the flow but no config path. Release Notes 2022 (2026-09-18)

32 **Background Tools remote command line and file transfer** NEW

45

Adds remote command line and remote file transfer capabilities via Background Tools, allowing admins to access and troubleshoot end-user devices without interrupting their work.

— States capability, no configuration steps. Release Notes 2023 (2026-09-18)

33 **Self-Service Account Provisioning** NEW

45

Introduces Self-Service Account Provisioning, letting new users bind their account to JumpCloud-managed macOS and Windows devices directly from the login window.

— Describes the flow briefly. Release Notes 2023 (2026-09-18)

34 **Portal UI refinements: password date, application favorites, quick search**

IMPROVED 45

Adds a Password Date column (last time a user's password was set) to the Users page in Admin Portal, splits the User Portal Applications area into All Applications and Favorites tabs with starring and responsive search, and adds a `Cmd/Ctrl + K` keyboard shortcut to the Admin Portal for quick navigation and search.

— Three named UI tweaks with thin mechanism. Release Notes 2024 (2026-09-18) Release Notes 2023 (2026-09-18)

35 **JumpCloud MCP Server (GA)** NEW 45

Introduces JumpCloud MCP Server (GA), connecting compatible AI tools and code editors to JumpCloud via the Model Context Protocol.

— Names the protocol, no usage detail. Release Notes 2025 (2026-09-18)

36 **IPA file support in Private Repository** NEW 45

Adds IPA file support to the JumpCloud Private Repository, enabling upload and deployment of IPA-based app packages to iOS and iPadOS devices.

— Names file type and target OS, no steps. Release Notes 2025 (2026-09-18)

37 **SaaS License Management (GA)** NEW 45

Adds SaaS License Management (GA) with automated license tracking and cost optimization via connectors.

— Names outcome, no mechanism detail. Release Notes 2025 (2026-09-18)

38 **Linux Forbidden Services Policy** NEW 40

Adds a Linux Forbidden Services Policy to remove unsupported service types from Linux fleets.

— Single policy, thin description. Release Notes 2022 (2026-09-18)

39 **Legacy Password Manager CSV export enhanced for Vault migration**

IMPROVED 40

Enhanced CSV export for the legacy Password Manager adds new data fields to support migration into JumpCloud Vault.

— States migration purpose without listing fields.

What's New: Password Manager Desktop Ap…

40 **ADI installation on Windows Server 2022** IMPROVED 40

Enables ADI installation on Windows Server 2022 domain controllers and member servers without workarounds or reduced functionality.

— States enablement, no mechanism detail. Release Notes 2023 (2026-09-18)

41 **Password reset from User Portal when MFA enforced** IMPROVED 40

Enables users to reset passwords from the User Portal when MFA is enforced but not yet enrolled.

— States behaviour, no walkthrough. Release Notes 2024 (2026-09-18)

42 **Service Account for APIs** NEW 40

Adds Service Account for APIs, allowing API service accounts not tied to individual Admin accounts.

— Bare capability statement. Release Notes 2025 (2026-09-18)

43 **AI Gateway integrations for Notion, Postman, Sentry, Zapier, Zoom**

NEW 35

The AI Gateway adds Notion, Postman, Sentry, Zapier, and Zoom as supported integrations.

— Names five integrations but gives no setup detail. product docs

44 **JumpCloud OpenID Connect (OIDC) SSO** NEW 35

Introduces JumpCloud OpenID Connect (OIDC) Single Sign-On (SSO) for OIDC-capable web applications.

— Names the protocol but no setup detail. Release Notes 2022 (2026-09-18)

45 **Live chat support for Premium Support customers** NEW 35

Adds live chat support for Premium Support customers directly within the Admin Portal, with automatic Premium Support status recognition (no PIN required).

— Simple support-channel addition, minimal mechanism. Release Notes 2022 (2026-09-18)

46 **WebAuthn enrollment split into Security Keys and Device authenticators**

IMPROVED 35

Adds WebAuthn enrollment split into separate credential types — Security Keys and Device authenticators — for clearer enrollment.

— Names two credential types only. Release Notes 2023 (2026-09-18)

47 **Source IP tracking for LDAP requests** NEW 35

Adds source IP tracking for incoming LDAP requests for enhanced security and traceability.

— One-line tracking addition. Release Notes 2025 (2026-09-18)

48 **Linux Fedora 37 agent support** NEW 30

Adds JumpCloud agent support for Linux Fedora 37.

— Bare platform-support note. Release Notes 2022 (2026-09-18)

49 **MFA enforced by default for new admin accounts** IMPROVED 25

Enables MFA by default for all new admin accounts.

— Bare default-change statement. Release Notes 2023 (2026-09-18)

└──▷ ALSO FROM THESE RELEASES

Stay current on JumpCloud product releases and agent updates by wiring the release notes feed into your RSS reader or SIEM ingest pipeline.

📍 Go to the JumpCloud Release Notes and Bug Fixes page › click Subscribe to Release Notes and Bug Fixes RSS Feed › copy the RSS URL into your RSS reader or feed aggregator.

└──▷ BREAKING ON UPGRADE

! The `GET` request for `api/systemusers/reactivate/{id}` is deprecated and now returns a 404 error; replace it with `POST` `https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups API `memberQuery` filter operators `$ge` and `$le` are deprecated; replace with `$gte` and `$lte` respectively in all API requests.

! LDAP no longer accepts connections using CBC-mode or SHA-1/SHA-256 integrity cipher suites with RSA or ECDHE key exchange — any LDAP client configured to use those suites will fail to connect.

! The `.appimage` Linux build is no longer supported as of v. 2.2.197 (April 18, 2024); users must migrate to the RPM or Debian build.

Release Notes 2022 (2026-09-18) BLOG ↗ STABLE

JumpCloud 2022 adds passwordless RADIUS auth, TOTP for LDAP, Remote Assist, Password Manager, and new MSP command templates.

└──▷ **HOW TO FIND IT**

Enforce a conditional access policy that blocks login from any device without full-disk encryption — useful for meeting compliance requirements across mixed OS fleets.

> 📍 In the Admin Portal, go to Security › Conditional Access, create or edit a policy, add a Device Condition, and set the encryption requirement to `FileVault enabled` (macOS), `system drive encrypted` (Windows), or 'root disk is encrypted' (Linux), then set the policy action to Deny.

Keep Google Chrome patched across every managed OS without creating per-platform policies — ideal for orgs running mixed macOS, Windows, and Linux fleets.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy, select `Universal Browser Patch Management`, choose Google Chrome as the target browser, configure the patch schedule, and assign the policy to the desired device group.

Run a consistent SentinelOne agent deployment across all managed endpoints without leaving the JumpCloud console.

> 📍 In the Admin Portal, go to Commands, create a new Command, select the target OS (macOS, Windows, or Linux), paste the SentinelOne agent download and install script, then run it against the target device group.

Adds `Command After Agent Install` command type so admins can specify enrollment commands to run immediately after the JumpCloud agent is installed on new devices.

Adds a `Running` column to the commands queue to distinguish commands waiting in the queue from commands that are actively running.

Introduces certificate-based (passwordless) authentication to JumpCloud RADIUS, supporting any X.509-compliant certificate authority — organizations can import existing certificates for RADIUS auth.

Adds TOTP as a supported MFA factor for LDAP applications, in addition to the existing JumpCloud Protect push-notification factor.

Adds a disk encryption device condition to Conditional Access Policies — admins can allow or deny access based on FileVault (macOS), system drive encryption (Windows), or root disk encryption (Linux) status.

+19 more

Release Notes 2023 (2026-09-18) BLOG ↗ STABLE

JumpCloud 2023 adds passwordless Go auth, Android zero-touch EMM, dynamic groups, remote background tools, and MSP billing visibility.

└──▷ **HOW TO FIND IT**

Automate device group membership so that newly enrolled macOS endpoints are grouped by vendor without manual assignment.

> 📍 In the Admin Portal, go to Device Groups › Create Group, enable Dynamic, add a rule for the Vendor attribute (e.g., Vendor equals `Apple`), and optionally enable `Review before applying` to preview membership changes before they take effect.

Protect admin accounts against MFA push bombing by tuning the concurrent attempt limit in MFA Configurations.

> 📍 In the Admin Portal, go to Security › MFA Configurations › JumpCloud Protect, locate Push Bombing Mitigation, and adjust the maximum concurrent attempts limit or disable the 60-second block window to match your organisation's risk tolerance.

Adds JumpCloud Go browser extension enabling hardware-protected, phishing-resistant passwordless login using Apple Touch ID and Windows Hello on managed macOS and Windows devices.

Enforces user and device verification on every SSO application session established via JumpCloud Go, with conditional access policies now defaulting to JumpCloud Go as the MFA method.

Supports Android zero-touch enrollment for fully managed, dedicated, and company-owned Android devices with a work profile.

Adds remote command line and remote file transfer capabilities via Background Tools, allowing admins to access and troubleshoot end-user devices without interrupting their work.

Introduces Self-Service Account Provisioning, letting new users bind their account to JumpCloud-managed macOS and Windows devices directly from the login window.

+24 more

Release Notes 2024 (2026-09-18) BLOG ↗ STABLE

JumpCloud 2024 adds Device Monitoring & Alerting, Windows App Catalog, SaaS Management, CSP-based MDM policies, and API filter improvements.

└──▷ **TRY IT**

Query users whose department is `Finance | Accounting` or `IT`, using the newly escapable pipe character in an `in` filter.

```
$ curl -X GET 'https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT' \
  -H 'x-api-key: <your-api-key>'
```

Enforce MFA for all MSP admins in bulk from the Multi-Tenant Portal to close gaps left by per-account settings.

> 📍 In the Admin Portal, go to Account › Settings and enable the `Global MFA Requirement` toggle under the Global MFA Requirement section.

Update an existing automation that reactivates user accounts to use the new POST endpoint before the deprecated GET starts returning 404.

```
$ curl -X POST
'https://console.jumpcloud.com/systemusers/reactivate/<user-id>' \
  -H 'x-api-key: <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d '{"activation_email": "user@example.com"}'
```

Adds `|` (pipe) escape support via `\` in `in` and `nin` API filter operators, enabling pipe characters within search terms — e.g. `GET /systemusers?filter=department:$in:Finance | Accounting|IT`.

Adds six new Windows MDM policies via CSP framework: Windows Wi-Fi Configuration Policy, Windows VPN Policy, Install Certificate Policy, Bluetooth Configuration Policy, Control Policy Conflict, and Windows Notification Service.

Adds a new Windows MDM policy to configure Simple Certificate Enrollment Protocol (`SCEP`) for Windows devices.

Introduces JumpCloud Device Monitoring and Alerting — monitor device health (availability, software, uptime, disk, battery), policy/command/software management results, directory events (user addition, privilege elevation), and custom alerts via command execution monitoring, manageable from the Admin Portal or MTP; supports Windows, macOS, and Linux.

Introduces the JumpCloud App Catalog for Windows, enabling install and auto-update management of 1Password, Adobe Acrobat Reader DC, Chrome, Firefox, Google Drive, OneDrive, OpenVPN, Slack, Webex, and Zoom.

+11 more

> └──▷ **BREAKING ON UPGRADE**

! The `GET` request for `api/systemusers/reactivate/{id}` is deprecated and now returns a 404 error; replace it with `POST` `https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups API `memberQuery` filter operators `$ge` and `$le` are deprecated; replace with `$gte` and `$lte` respectively in all API requests.

`Release Notes 2025 (2026-09-18)` BLOG ↗ STABLE

JumpCloud: Adds Authentication Methods References (AMR) claims to SAML assertions and OIDC ID tokens, tagging logins with proof of factors used (e.g., password + OTP).

└──▷ **HOW TO FIND IT**

Scope a procurement team member's admin access to Asset Management only, without granting broader Admin Portal rights.

📍 In the Admin Portal, go to Administrators › Roles, assign the built-in `Asset Manager` role to the admin account for the procurement team member.

Adds support for `ForceAuthn="true"` in SAML and `prompt=login` in OIDC so applications can trigger step-up authentication without ending a user's existing session.

Introduces JumpCloud Access Requests, enabling configurable approval flows, in-console request management, Directory Insights event tracking, and webhook integration with ticketing systems such as JIRA.

Introduces JumpCloud MCP Server (GA), connecting compatible AI tools and code editors to JumpCloud via the Model Context Protocol.

Introduces AI Commands Builder (GA) to generate and refine commands for Windows, macOS, and Linux devices from plain-English prompts inside the Admin Portal.

Introduces the Shadow AI Dashboard, giving visibility into generative AI app usage, adoption by department, and most active users within JumpCloud AI & SaaS Management.

`+14 more`

└──▷ **BREAKING ON UPGRADE**

! LDAP no longer accepts connections using CBC-mode or SHA-1/SHA-256 integrity cipher suites with RSA or ECDHE key exchange — any LDAP client configured to use those suites will fail to connect.

`Release Notes 2026 (2026-09-18)` BLOG ↗ STABLE

JumpCloud adds Password Vault across Admin/User portals and mobile, Agent Identities for non-human IAM, and GA Access Risk Detection with behavior-based anomaly alerting.

└──▷ HOW TO FIND IT

Register a new AI agent and scope its access — useful when onboarding an Amazon Bedrock AgentCore integration that requires `Client secret post` auth.

> 📍 1. Go to Identity Management › Agents and click + Add. 2. Enter the agent name, owner, and company email on the Details tab. 3. On the Authentication tab, generate OAuth client credentials and set the client authentication type to `Client secret post`. 4. Under Identity Management › Agent Groups, create an Agent Group, add the agent, and assign the target SSO applications. 5. On the agent Applications tab, click Manage connections to authorize the relevant MCP servers on behalf of the agent.

Surface and connect a device-discovered MCP server to AI Gateway from the SaaS Management view, without switching to a separate tool.

> 📍 1. Go to Access › AI & SaaS Management › Applications. 2. Check the MCP Connection column for an app showing `Discovered`. 3. Click the status to review the discovered MCP servers. 4. Click Add MCP Server to register it in AI Gateway.

Enable Access Risk Detection and route high-severity alerts to Slack and a webhook for SOC triage.

> 📍 1. Go to Configure Access Risk Detection in the Admin Portal. 2. Select the risk factors to monitor (e.g., impossible travel, unfamiliar devices, dormant accounts) and tune thresholds for your environment. 3. Under notifications, add your Slack channel and webhook URL, mapped to the appropriate severity levels. 4. Save — risky logins will now trigger alerts and auto-resolve when the user completes MFA.

Adds `Access > AI & SaaS Management > Applications` view with an `MCP Connection` column showing `Discovered`, `Connected`, or `Not Available` status for each app, plus an `Add MCP Server` action to connect discovered MCP servers to AI Gateway without switching tools.

Adds configurable client authentication type (`Client secret basic` or `Client secret post`) on the agent `Authentication` tab for OAuth credential requests, enabling compatibility with platforms such as Amazon Bedrock AgentCore.

Adds Password Vault to the Admin Portal, enabling administrators to enable the vault, assign access to users and user groups, manage shared credentials and websites with `Connect`, `View`, `Edit`, and `Manage` permissions, monitor password health from an Overview dashboard, and view audit logs in Directory Insights.

Adds Password Vault to the User Portal (left navigation after login), supporting credential and website management, shared folders, and bulk import via CSV or Excel file.

Adds Password Vault autofill to the JumpCloud Go browser extension, supporting non-SSO website password autofill, 2FA code launch, secret viewing, and save-on-login/update/signup without device registration.

+1 more

What's New: Password Manager Desktop App (2026-09-18) BLOG ↗ LATEST STABLE

JumpCloud Password Manager Desktop App adds EU region support, 30K-item capacity, ARM build for Apple Silicon, and admin shared-folder management.

Adds region selection on the activation page (desktop app) and pairing page (browser extension), with EU region now supported.

Scales item storage to up to 30,000 items (without autofill) or 13,000 items (with autofill); a single shared folder can hold up to 5,000 items.

Introduces a Move All items button in shared folders to relocate all items from a folder in one action.

Introduces an auto-lock schedule with a default of 15 minutes of inactivity and a maximum of 1,440 minutes (24 hours).

Adds an ARM build for macOS devices with M1 and above silicon, available on the Downloads page, with automatic upgrade from x64.

+11 more

└──▷ BREAKING ON UPGRADE

! The `.appimage` Linux build is no longer supported as of v. 2.2.197 (April 18, 2024); users must migrate to the RPM or Debian build.

announced · Sep 15, 2026 BLOG ↗

JumpCloud Password Vault is now GA: a centralized, encrypted enterprise credential manager with audit logs, health dashboard, and mobile MFA integration.

└──▷ HOW TO FIND IT

Review the Password Health Dashboard to proactively identify weak or expired credentials across the organization before a breach forces a reactive response.

📍 In the JumpCloud Admin Portal, go to Password Vault › Password Health Dashboard to view organization-wide credential health, including expired, expiring, weak, and inactive credentials.

Reassign a departing employee's credentials to a colleague to prevent business disruption during offboarding.

> 📍 In the JumpCloud Admin Portal, navigate to Password Vault, locate the orphaned folder created for the offboarded user, and assign it to the replacement user.

Adds JumpCloud Password Vault, a new cloud-first enterprise credential management solution that centralizes all company passwords and secrets in a single encrypted vault with unlimited organizational credential storage.

Adds a Password Health Dashboard providing organization-wide visibility into expired, expiring, weak, and inactive credentials alongside a built-in password generator to enforce strong, unique passwords.

Adds four distinct credential sharing permission levels — Connect, View Secret, View Details, and Manage — replacing ad-hoc sharing over email or chat.

Extends the JumpCloud Go browser extension to unify Password Vault management, JumpCloud Go authentication, SSO access, and SaaS management in a single extension, enabling automatic sign-in to non-SSO apps without exposing the password.

Extends JumpCloud Protect mobile app to combine corporate MFA and Password Vault access in a single app, eliminating the need for separate authenticator and password manager apps.

+2 more

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ

Custom User ID Assignment with Actions — Early Access (2026-09-17) BLOG ↗

RANK

Secure users, AI agents, and more with Auth0, an easy-to-implement, scalable, and adaptable authentication and authorization platform.

> Auth0 shipped early-access custom user_id assignment via Actions, a new My Organization API with bulk invitation and role-based member management (plus embeddable UI components), and a native v0 integration that wires real Auth0 tenants into generated apps.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **Member management in the My Organization API** NEW 88

Early Access adds bulk member invitation (up to 10 invitees in a single atomic request, with pre-assigned roles and routing through a specific identity provider or user store), member listing/retrieval with field projection and checkpoint-based pagination, and role assignment/removal through the API or embeddable UI components — with guardrails preventing delegated admins from assigning roles exceeding their own. A new `none`, `readonly`, `limited`, and `full` access level model is computed automatically from a member's organization connection access levels, embeddable UI components show or hide actions based on effective permissions sourced from the ID token with no extra API call, and tenant log events are emitted for all member, invitation, and role operations. Enable it via `Applications > APIs` in the Auth0 Dashboard.

Activate the My Organization API for your tenant so organization administrators can begin managing members, invitations, and roles.

> 📍 In the Auth0 Dashboard, go to Applications › APIs and activate the My Organization API.

— Names limits, access-level model, log events and exact dashboard activation path.

My Organization API and Embeddable UI C…

02 **Custom user_id assignment via pre-user-registration Actions**

NEW 80

Adds api.user.setUserId(), callable inside a `pre-user-registration` Action, to assign a custom `user_id` — UUIDs, domain-specific IDs, or legacy-compatible IDs — to new database-connection users at signup, replacing the automatic `auth0|<id>` format. Custom user IDs are supported across Database connections (including custom

databases with automatic migration), Universal Login sign-up, the Authentication API signup endpoint, and the Management API Create User endpoint.

Assign a UUID-format user_id at signup to keep identity consistent between newly-registered users and records imported from a legacy system.

```javascript
exports.onExecutePreUserRegistration = async (event, api) => {
  const crypto = require('crypto');
  api.user.setUserId(crypto.randomUUID());
};
```

— Names exact method, hook, and four supported surfaces with runnable code.

Custom User ID Assignment with Actions …

03 **Native Auth0 integration for v0** NEW 72

A new Auth0 integration for v0 lets prompts create Auth0 applications, set up identity connections, and wire the Auth0 SDK into generated code without leaving v0's conversational workflow, supporting B2B use cases, enterprise identity provider connections, and MFA requirements configured directly from prompts against a real Auth0 tenant. Install it from v0 project settings under the `Integrations` tab by searching for Auth0 and clicking `Install`, then linking or creating an Auth0 account.



Add real Auth0 authentication — including protected routes and session handling — to a v0-built app without leaving your prompt session.

> 📍 In v0, open your project settings, navigate to the `Integrations` tab, search for Auth0, click `Install`, and complete the account link. Then prompt: 'Allow users to sign up and log in to access the application using the Auth0 integration.'

— Exact UI path and prompt example given, but mechanism stays high-level.

launch-20260916-0bc01060

└──▷ ALSO FROM THESE RELEASES

Custom User ID Assignment with Actions — Early Access (2026-09-17) BLOG ↗

STABLE

Auth0 Actions can now assign custom user_id values at signup via api.user.setUserId() in pre-user-registration hooks.

└──▷ USE IT

Assign a UUID-format user_id at signup to keep identity consistent between newly-registered users and records imported from a legacy system.

javascript

```
exports.onExecutePreUserRegistration = async (event, api) => {
  const crypto = require('crypto');
  api.user.setUserId(crypto.randomUUID());
};
```

Adds api.user.setUserId() method callable inside a `pre-user-registration` Action to assign a custom `user_id` — UUIDs, domain-specific IDs, or legacy-compatible IDs — to new database-connection users at signup, replacing the automatic `auth0|<id>` format.

Supports custom user IDs across Database connections (including custom databases with automatic migration), Universal Login sign-up, the Authentication API signup endpoint, and the Management API Create User endpoint.

My Organization API and Embeddable UI Components - Member Management in Early Access!

BLOG ↗ LATEST STABLE

Auth0 My Organization API adds Member Management in Early Access: bulk invitations, role assignment, and embeddable UI components for B2B SaaS.

└──▷ HOW TO FIND IT

Activate the My Organization API for your tenant so organization administrators can begin managing members, invitations, and roles.

📍 In the Auth0 Dashboard, go to Applications › APIs and activate the My Organization API.

Adds bulk member invitation via the My Organization API, supporting up to 10 invitees in a single atomic request with pre-assigned roles and routing through a specific identity provider or user store.

Adds member listing and retrieval through the My Organization API with field projection and checkpoint-based pagination for large organizations.

Adds role assignment and removal for organization members directly through the My Organization API or embeddable UI components, with permission guardrails preventing delegated administrators from assigning roles exceeding their own.

Introduces a new `none`, `readonly`, `limited`, and `full` access level model for organization members, computed automatically from the member's organization connection access levels.

Adds permission-aware embeddable UI components that show or hide actions based on the authenticated user's effective permissions sourced from the ID token, requiring no additional API call.

+2 more

announced · Sep 16, 2026 BLOG ↗

Auth0 launches a native integration for v0 that wires real Auth0 tenants — apps, connections, and SDKs — directly from prompts.

└──▷ **HOW TO FIND IT**

Add real Auth0 authentication — including protected routes and session handling — to a v0-built app without leaving your prompt session.

> 📍 In v0, open your project settings, navigate to the `Integrations` tab, search for Auth0, click `Install`, and complete the account link. Then prompt: 'Allow users to sign up and log in to access the application using the Auth0 integration.'

New Auth0 integration for v0 lets prompts create Auth0 applications, set up identity connections, and wire the Auth0 SDK into generated code without leaving the v0 conversational workflow.

Supports B2B use cases, enterprise identity provider connections, and MFA requirements configured directly from v0 prompts against a real Auth0 tenant.

Install via v0 project settings under the `Integrations` tab by searching for Auth0 and clicking `Install`, then follow on-screen steps to link or create an Auth0 account.

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Weekly release 2026.09.1 (2026-09-16) BLOG ↗ RANK

Secure, scalable, and highly available authentication and user management for any app.

> Okta Identity Engine expanded AI-agent support with new Beta APIs and broader CIMD coverage, added scoped token exchange and a GA user identification policy, while Identity Governance introduced EA staging for certification campaigns before launch.

└──▷ **WHAT SHIPPED · 5 FEATURES** most completely described first ? what's the number?

01 **Staged certification campaign preview (EA)** NEW 55

Adds Early Access support for staging certification campaigns before launch, letting admins preview target users, resources, and reviewer assignments to validate scope and fix configuration errors prior to go-live.

– Explains preview mechanism but no UI path or flag named product docs

02 **CIMD expansion to AI Agents and OIDC apps** NEW 50

Continuous Identity and Metadata Delivery (CIMD) now extends to AI Agents (Beta), covering agent-based workflows, and CIMD for OIDC apps is now Generally Available in Preview for OpenID Connect applications.

– Names the two CIMD targets and stages but no config or endpoint given

Weekly release 2026.09.1 (2026-09-16)

03 **Beta APIs for AI agent traffic and MCP server discovery** NEW 50

Introduces Agent Gateway APIs (Beta) for managing AI agent traffic through Okta, and MCP server discovery APIs (Beta) to locate and integrate Model Context Protocol servers.

– Names both new API surfaces but no endpoint paths given

Weekly release 2026.09.1 (2026-09-16)

04 **Scoped token exchange requests** IMPROVED 35

Token exchange requests now support scopes, enabling finer-grained access control when exchanging tokens.

– Describes the change but no scope names or request example

Weekly release 2026.09.1 (2026-09-16)

05 **User identification policy GA in Preview** NEW 35

User identification policy is now GA in Preview, providing a dedicated policy layer for identifying users during authentication flows.

— States the capability but no configuration detail Weekly release 2026.09.1 (2026-09-16)

Weekly release 2026.09.1 (2026-09-16) BLOG ↗ LATEST STABLE

Okta adds CIMD for AI Agents, Agent Gateway APIs, MCP server discovery, and token exchange scope support in 2026.09.1

Token exchange requests now support scopes, enabling finer-grained access control when exchanging tokens.

Introduces CIMD for AI Agents (Beta), extending Continuous Identity and Metadata Delivery to agent-based workflows.

Introduces Agent Gateway APIs (Beta) for managing AI agent traffic through Okta.

Introduces MCP server discovery APIs (Beta) to locate and integrate Model Context Protocol servers.

CIMD for OIDC apps is now GA in Preview, making Continuous Identity and Metadata Delivery generally available for OpenID Connect applications.

+1 more

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ announced · Sep 15, 2026 BLOG ↗

RANK

Govern workforce access, deploy AI agents, and get work done with C1.

> ConductorOne launched an EU-hosted instance to help European customers meet data-residency requirements.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **EU-hosted instance in AWS Frankfurt** NEW 50

C1.ai now offers a new EU instance that provisions tenant data in AWS Frankfurt, with Ireland serving as the disaster-recovery region, enabling European customers to meet data-residency requirements.

— Names region and DR site but no migration or setup steps launch-20260915-39e4b833

> announced · Sep 15, 2026 BLOG ↗
>
> ConductorOne is now available as an EU-hosted instance, with tenant data held in AWS Frankfurt.
>
> New EU instance of C1.ai provisions tenant data in AWS Frankfurt (Ireland as disaster-recovery region), enabling European customers to meet data-residency requirements.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Firezone

Feedback ⓘ announced · Sep 14, 2026 BLOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone launched an official Terraform provider that brings its access model into infrastructure-as-code, covering Sites, Resources, Policies and Gateways, with fine-grained resource filtering, condition-based policies, and gateway token rotation.

└──▷ WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 **firezone_policy condition-based access rules** NEW 95

`firezone_policy` supports `condition` blocks that gate access on `auth_provider_id`, `current_utc_datetime` (day-of-week/time-range windows in any IANA timezone), `device_attested`, device region, and remote IP (CIDR) match properties, letting policies like business-hours-only or attested-device-only access be defined and reviewed as code.

Lock down a sensitive resource to Okta-authenticated Engineering members during business hours on attested devices — defined entirely in code so the policy is reviewable in a PR.

```hcl
resource "firezone_resource" "database" {
  site_id = firezone_site.main.id
  name    = "postgres-prod"
  type    = "ip"
  address = "10.0.0.100"
  filters {
    protocol = "tcp"
    ports    = ["5432"]
  }
}

resource "firezone_policy" "engineering_db" {
  group_id    = data.firezone_group.engineering.id
  resource_id = firezone_resource.database.id
  condition {
    property = "auth_provider_id"
    operator = "is_in"
    values   = [data.firezone_okta_auth_provider.corp.id]
  }
  condition {
    property = "current_utc_datetime"
    operator = "is_in_day_of_week_time_ranges"
    values   = ["M/09:00-17:00/America/New_York", "F/09:00-
17:00/America/New_York"]
  }
  condition {
    property = "device_attested"
    operator = "is"
    values   = ["true"]
  }
}
```

— Names every condition property with mechanism and a full runnable policy example

launch-20260914-a772d686

02 **Official Terraform provider for Firezone IaC** NEW 85

Adds the official `firezone/firezone` Terraform provider, published on the Terraform Registry, enabling `firezone_site`, `firezone_resource`, `firezone_policy`, and `firezone_gateway` resources to be managed as code. Authentication is done via the `FIREZONE_TOKEN` environment variable and an `endpoint` provider config key (defaulting to `https://rest-api.firezone.dev`), using an API client token created in the Portal under Settings.

Bootstrap the provider using an API client token from the Firezone Portal so all subsequent `terraform apply` runs authenticate automatically.

```hcl
terraform {
  required_providers {
    firezone = {
      source = "firezone/firezone"
    }
  }
}

provider "firezone" {
  endpoint = "https://rest-api.firezone.dev"
}

# Set FIREZONE_TOKEN in your environment:
# export FIREZONE_TOKEN=<token-from-portal-settings>
```

— Names provider source, env var, config key and default endpoint with runnable example

launch-20260914-a772d686

---

03 **firezone_gateway token rotation trigger** NEW 75

`firezone_gateway` adds a `token_rotation_trigger` attribute that rotates the Gateway's enrollment token whenever its value changes, storing the sensitive token in Terraform state; the old token remains valid until the Gateway reconnects with the new one.

Rotate a Gateway token on a schedule by updating `token_rotation_trigger` — the old token stays valid until the Gateway reconnects with the new one.

```hcl
resource "firezone_gateway" "gw" {
  for_each = {
    "a" = { name = "gw-us-east-1a", rotated = "2026-09-14" }
  }
  site_id              = firezone_site.main.id
  name                 = each.value.name
  token_rotation_trigger = each.value.rotated
}
```

— Names the attribute and explains rotation behavior with an example

launch-20260914-a772d686

---

04 **firezone_resource types with traffic filters** NEW 70

`firezone_resource` supports `cidr`, `ip`, and `dns` resource types, with `filters` blocks to restrict by protocol and port, and an `ip_stack` attribute for DNS resources.

Lock down a sensitive resource to Okta-authenticated Engineering members during business hours on attested devices — defined entirely in code so the policy is reviewable in a PR.

hcl

```
resource "firezone_resource" "database" {
  site_id = firezone_site.main.id
  name    = "postgres-prod"
  type    = "ip"
  address = "10.0.0.100"
  filters {
    protocol = "tcp"
    ports    = ["5432"]
  }
}

resource "firezone_policy" "engineering_db" {
  group_id    = data.firezone_group.engineering.id
  resource_id = firezone_resource.database.id
  condition {
    property = "auth_provider_id"
    operator = "is_in"
    values   = [data.firezone_okta_auth_provider.corp.id]
  }
  condition {
    property = "current_utc_datetime"
    operator = "is_in_day_of_week_time_ranges"
    values   = ["M/09:00-17:00/America/New_York", "F/09:00-
17:00/America/New_York"]
  }
  condition {
    property = "device_attested"
    operator = "is"
    values   = ["true"]
  }
}
```

— Names resource types and filter attributes with a working config example

launch-20260914-a772d686

THINNER COVERAGE BELOW

05 **Group and Okta auth provider data sources** NEW 40

Adds `firezone_group` and `firezone_okta_auth_provider` Terraform data sources for referencing existing groups and auth provider integrations when building Policies.

Lock down a sensitive resource to Okta-authenticated Engineering members during business hours on attested devices — defined entirely in code so the policy is reviewable in a PR.

hcl

```
resource "firezone_resource" "database" {
  site_id = firezone_site.main.id
  name    = "postgres-prod"
  type    = "ip"
  address = "10.0.0.100"
  filters {
    protocol = "tcp"
    ports    = ["5432"]
  }
}

resource "firezone_policy" "engineering_db" {
  group_id    = data.firezone_group.engineering.id
  resource_id = firezone_resource.database.id
  condition {
    property = "auth_provider_id"
    operator = "is_in"
    values   = [data.firezone_okta_auth_provider.corp.id]
  }
  condition {
    property = "current_utc_datetime"
    operator = "is_in_day_of_week_time_ranges"
    values   = ["M/09:00-17:00/America/New_York", "F/09:00-
17:00/America/New_York"]
  }
  condition {
    property = "device_attested"
    operator = "is"
    values   = ["true"]
  }
}
```

— Names two data sources but gives no mechanism beyond referencing

launch-20260914-a772d686

announced · Sep 14, 2026 BLOG ↗

Firezone launches an official Terraform provider, enabling infrastructure-as-code management of Sites, Resources, Policies, and Gateways.

└──▷ USE IT

Bootstrap the provider using an API client token from the Firezone Portal so all subsequent `terraform apply` runs authenticate automatically.

main.tf

```
terraform {
  required_providers {
    firezone = {
      source = "firezone/firezone"
    }
  }
}

provider "firezone" {
  endpoint = "https://rest-api.firezone.dev"
}

# Set FIREZONE_TOKEN in your environment:
# export FIREZONE_TOKEN=<token-from-portal-settings>
```

Lock down a sensitive resource to Okta-authenticated Engineering members during business hours on attested devices — defined entirely in code so the policy is reviewable in a PR.

main.tf

```
resource "firezone_resource" "database" {
  site_id = firezone_site.main.id
  name    = "postgres-prod"
  type    = "ip"
  address = "10.0.0.100"
  filters {
    protocol = "tcp"
    ports    = ["5432"]
  }
}

resource "firezone_policy" "engineering_db" {
  group_id    = data.firezone_group.engineering.id
  resource_id = firezone_resource.database.id
  condition {
    property = "auth_provider_id"
    operator = "is_in"
    values   = [data.firezone_okta_auth_provider.corp.id]
  }
  condition {
    property = "current_utc_datetime"
    operator = "is_in_day_of_week_time_ranges"
    values   = ["M/09:00-17:00/America/New_York", "F/09:00-
17:00/America/New_York"]
  }
  condition {
    property = "device_attested"
    operator = "is"
    values   = ["true"]
  }
}
```

Rotate a Gateway token on a schedule by updating `token_rotation_trigger` — the old token stays valid until the Gateway reconnects with the new one.

main.tf

```
resource "firezone_gateway" "gw" {
  for_each = {
    "a" = { name = "gw-us-east-1a", rotated = "2026-09-14" }
  }
  site_id              = firezone_site.main.id
  name                 = each.value.name
  token_rotation_trigger = each.value.rotated
}
```

Adds the official `firezone/firezone` Terraform provider, available on the Terraform Registry, enabling Sites, Resources, Policies, and Gateways to be managed as code.

Adds `FIREZONE_TOKEN` environment variable and `endpoint` provider config key (defaulting to `https://rest-api.firezone.dev`) for authenticating the Terraform provider via an API client token created in the Portal under Settings.

Adds `firezone_site`, `firezone_resource`, `firezone_policy`, and `firezone_gateway` Terraform resources for managing the core Firezone access model.

Adds `firezone_resource` with support for `cidr`, `ip`, and `dns` resource types, with `filters` blocks for protocol and port restrictions, and `ip_stack` for DNS resources.

Adds `firezone_policy` with `condition` blocks supporting `auth_provider_id`, `current_utc_datetime` (day-of-week/time-range windows in any IANA timezone), `device_attested`, device region, and remote IP (CIDR) match properties.

+2 more

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.47.10 (2026-09-17) BLOG ↗ RANK

RegScale is a compliance and security governance platform that automates regulatory assessment, documentation, and continuous monitoring across organizations.

> RegScale CLI 6.47.10 introduces shared security-plan lookup infrastructure and uses it to add multi-plan routing commands for QRadar events and AWS Security Hub CSPM findings.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **Multi-plan routing for QRadar and AWS Security Hub CSPM** NEW

60

Adds a shared lookup that finds a security plan by an identifier held in any of its fields, native or custom. Built on this, a new QRadar command routes events to many security plans by an identifier carried on the event, and a new AWS Security Hub CSPM command routes findings to many security plans.

— Names the commands and lookup mechanism but no exact CLI syntax or flags

CLI 6.47.10 (2026-09-17)

> CLI 6.47.10 (2026-09-17) BLOG ↗ LATEST STABLE
>
> RegScale CLI 6.47.10 adds shared security-plan lookup, multi-plan QRadar routing, and AWS Security Hub CSPM command
>
> Adds a shared lookup that finds a security plan by an identifier held in any of its fields, native or custom.
>
> Adds a QRadar command that routes events to many security plans by an identifier carried on the event.
>
> Adds an AWS Security Hub CSPM command that routes findings to many security plans.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ

1.28.545 - Slack thread following, on-demand schedules, and trigger management over MCP

BLOG ↗ RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

> Runlayer shipped a full Microsoft Teams integration this window — per-agent bots, ambient channel replies, outbound posting, workspace-wide notifications, and resource link previews — alongside Slack thread continuity, on-demand schedule execution, MCP-based trigger management, new default models, and a breaking change to the agent channel intro message.

└──▷ WHAT SHIPPED · 10 FEATURES most completely described first ? what's the number?

01 Workspace-level Teams notifications and security alerts NEW

85

A shared Runlayer Teams Bot provides workspace-level notifications, configured from Settings → Integrations → Microsoft Teams with a Notification channel field, an optional Security alerts channel for security scan violations, and an Only alert on blocked calls toggle to limit alerts to blocked calls.

Route security scan violation alerts to a dedicated SOC channel and suppress low-severity mask/alert-only noise.

📍 1. Go to Settings → Integrations → Microsoft Teams. 2. Click Create Teams Bot and complete the tenant connection if not already done. 3. @mention the bot at least once from your #soc-alerts channel so it can post there. 4. Set Security alerts channel to #soc-alerts. 5. Enable Only alert on blocked calls. 6. Click Test next to the Security alerts channel to confirm delivery.

— Names every config key and includes full setup steps. product docs

02 Ambient channel replies without @mention NEW

75

A Respond without being tagged in channels toggle under each agent's Microsoft Teams settings lets the bot answer channel messages it isn't @mentioned in, with an optional natural-language condition to scope which messages trigger a response.

Enable ambient channel monitoring so a security-focused agent answers every deployment-failure discussion without needing an @mention.

📍 1. Open the agent detail page. 2. Under Microsoft Teams, toggle on `Respond without being tagged in channels`. 3. In the condition field that appears, enter: only respond to messages about deployment failures 4. Click Save. 5. Reinstall the Teams app (Microsoft Teams → Install, same dialog) so the team owner can grant the extra permission.

— Exact toggle name plus a worked configuration example. product docs

03 **Per-agent Microsoft Teams bots and installation** NEW

65

Runlayer adds full Microsoft Teams integration, giving each agent its own bot that responds to @mentions and direct messages in Teams channels; org-wide tenant authorization is granted via `Settings → Integrations → Connect Microsoft Teams`. Each agent's bot can be installed from the agent detail page under `Microsoft Teams → Install`, targeting a specific user (direct message), a team (all channels), or bulk-installed for all users or all teams.

Install an agent bot across an entire team in bulk so every channel member can @mention it immediately.

📍 1. Open the agent detail page → Microsoft Teams → Install. 2. In the Install Teams app dialog, choose Team (all channels). 3. Click Install for all teams. 4. Click Install.

— Exact menu paths given but no API for automated setup. product docs

04 **Teams link previews for Runlayer resources** NEW

65

Pasting an agent, skill, plugin, artifact, or session URL into a Teams chat triggers the agent's bot to resolve it into a rich preview card showing resource type, title, and creator, matching Slack's existing link-preview behavior and respecting access controls; this requires reinstalling the Teams app to pick up the updated manifest.

— Names resource types and the reinstall step needed to enable it. product docs

05 **On-demand schedule execution** NEW

60

Schedules can now be run immediately or have failed runs retried directly from the schedule list, without waiting for the next scheduled window.

Immediately re-run a failed scheduled job without waiting for the next scheduled window.

📍 In the Runlayer console, go to Agents › your agent › Schedules, then click `Run now` on the target schedule or `Retry` next to a failed run.

— Exact navigation path and buttons (Run now/Retry) given.

1.28.545 - Slack thread following, on-d…

THINNER COVERAGE BELOW

06 **Slack thread continuity and attachment delivery** IMPROVED

50

Agents can now follow a Slack thread continuously after an initial mention without requiring re-mention for each reply, and Office document attachments shared in a Slack thread are delivered to the agent sandbox, with undeliverable files surfaced explicitly instead of silently skipped.

— Behavior change described but no toggle or config surface named.

1.28.545 - Slack thread following, on-d…

07 **Agent-initiated outbound posting to Teams channels** NEW

45

Agents can post a new message into any Teams channel whose ID is provided, independent of any incoming message that triggered it, enabling proactive outbound notifications from an agent.

— Capability named but no exact API call or command shown. product docs

08 **Agent channel intro message off by default** BREAKING

45

The agent channel intro message is now off by default; existing agents that relied on it broadcasting on channel join will no longer do so unless re-enabled per agent.

— Clear before/after and migration note but no exact toggle location.

1.28.545 - Slack thread following, on-d…

09 **Trigger management via MCP tools** NEW

40

Adds trigger management through Runlayer MCP tools — create, list, and delete agent triggers programmatically.

— No specific MCP tool names or call schema given.

1.28.545 - Slack thread following, on-d…

10 **Claude Sonnet 5 default model for new agents** IMPROVED

40

Sets Claude Sonnet 5 as the default model for newly created agents; the managed assistant itself now runs on Claude Opus 4.8.

— States new defaults but no override or migration instructions.

1.28.545 - Slack thread following, on-d…

└──▷ BREAKING ON UPGRADE

! The agent channel intro message is now off by default; existing agents that relied on it broadcasting on channel join will no longer do so unless re-enabled per agent.

`1.28.545 - Slack thread following, on-demand schedules, and trigger management over MC`

BLOG ↗ LATEST STABLE

Runlayer 1.28.545 adds Slack thread following, on-demand schedule runs, and trigger management via MCP tools.

└──▷ **HOW TO FIND IT**

Immediately re-run a failed scheduled job without waiting for the next scheduled window.

📍 In the Runlayer console, go to Agents › your agent › Schedules, then click `Run now` on the target schedule or `Retry` next to a failed run.

Adds trigger management via Runlayer MCP tools — create, list, and delete agent triggers programmatically.

Enables agents to follow a Slack thread continuously after an initial mention, without requiring re-mention for each reply.

Adds on-demand schedule execution — run a schedule immediately or retry failed scheduled runs directly from the schedule list.

Delivers Office document attachments shared in a Slack thread to the agent sandbox, and surfaces undeliverable files explicitly instead of silently skipping them.

Sets Claude Sonnet 5 as the default model for new agents; the managed assistant now runs on Claude Opus 4.8.

`+1 more`

└──▷ **BREAKING ON UPGRADE**

! The agent channel intro message is now off by default; existing agents that relied on it broadcasting on channel join will no longer do so unless re-enabled per agent.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ai-safe2-framework

Feedback ⓘ

2026-09-17_CLI_0.6.0_Scope_Attribute_Decide NOTES ↗ RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> AI SAFE² CLI 0.6.0 introduces a suite of `safe2 evidence` subcommands for scoped assessment, change attribution, and release-readiness reporting, backed by new exportable schema contracts and hardened JUnit XML intake.

└──▷ **WHAT SHIPPED · 5 FEATURES** most completely described first ? what's the number?

01 Evidence readiness command for release-readiness cards NEW

95

`safe2 evidence readiness` produces canonical release-readiness JSON and an escaped Markdown card, accepting `--system-identity`, `--assessment-scope`, `--change-attribution`, `--output`, `--card`, and `--strict` options. Every readiness output includes a permanent `release_authorized: false` field, enforcing that no automated result can be mistaken for release permission.

— Names command, all flags, and the permanent safety field

2026-09-17_CLI_0.6.0_Scope_Attribute_De…

02 Evidence scope command for deployment inventory NEW

90

`safe2 evidence scope` inventories the declared deployment relationship of repository or built-artifact paths, classifying each as `included`, `excluded`, `partial`, `not applicable`, or `unknown`.

— Names exact command and its classification categories

2026-09-17_CLI_0.6.0_Scope_Attribute_De…

03 Evidence attribute command for change attribution NEW

90

`safe2 evidence attribute` compares normalized findings across an explicitly trusted baseline and the current revision, separating them into inherited, introduced, changed, resolved, and unknown categories.

— Names exact command and its comparison categories

2026-09-17_CLI_0.6.0_Scope_Attribute_De…

04 New schema contracts for scope, attribution and readiness

NEW

75

Adds three new schema contracts — `assessment-scope-source-v1`, `change-attribution-source-v1`, and `release-readiness-source-v1` — each exportable via `safe2 schema export`.

— Names all three schema ids and export command, no further detail

2026-09-17_CLI_0.6.0_Scope_Attribute_De…

05 **Hardened JUnit XML intake against XML attacks** IMPROVED

65

JUnit XML intake now runs through `defusedxml`, rejecting DTD/entity input while retaining existing size, encoding, structure, and declared-total checks.

— Names the library and checks but no user-facing command

2026-09-17_CLI_0.6.0_Scope_Attribute_De…

2026-09-17_CLI_0.6.0_Scope_Attribute_Decide NOTES ↗ LATEST STABLE

AI SAFE² CLI 0.6.0 adds `safe2 evidence scope`, `attribute`, and `readiness` commands for scoped assessment, change attribution, and release-readiness cards.

└──▷ GET THIS VERSION

```
$ git clone --branch 2026-09-17_CLI_0.6.0_Scope_Attribute_Decide
https://github.com/CyberStrategyInstitute/ai-safe2-framework.git
# already have the repo? check out this version:
$ git checkout 2026-09-17_CLI_0.6.0_Scope_Attribute_Decide
```

Adds `safe2 evidence scope` command to inventory the declared deployment relationship of repository or built-artifact paths, classifying each as included, excluded, partial, not applicable, or unknown.

Adds `safe2 evidence attribute` command to compare normalized findings across an explicitly trusted baseline and current revision, separating inherited, introduced, changed, resolved, and unknown findings.

Adds `safe2 evidence readiness` command to produce canonical release-readiness JSON and an escaped Markdown card, accepting `--system-identity`, `--assessment-scope`, `--change-attribution`, `--output`, `--card`, and `--strict` options.

Adds three new schema contracts — `assessment-scope-source-v1`, `change-attribution-source-v1`, and `release-readiness-source-v1` — exportable via `safe2 schema export`.

Introduces `release_authorized: false` as a permanent field in every readiness output, enforcing that no automated result can be mistaken for release permission.

> +1 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 18, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue