THE CYBER TOOLCHAIN

NO. 084

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 24, 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 279 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 50 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

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. OpenBao v2.7.0 moves secrets infrastructure to post-quantum TLS and ML-DSA signing while letting

vim ? help

PKI and Transit keys stay in an external KMS. Wireshark can now attribute each packet to the process that sent it. And Elastic's Cross-Project Search lets one project triage alerts and hunt across up to 100 linked serverless projects.

PROTECT

Keep the secrets backend usable after the crypto changes underneath it

OpenBao v2.7.0 adds PQC TLS key exchange and ML-DSA signing, plus External Keys so PKI and Transit operations can run against a KMS-held key that never lands in Bao's storage. DigiCert is moving default TLS issuance to G5 roots and adds a persistent DNS TXT record that satisfies recurring revalidation, so domain control checks stop being a recurring manual task as validation reuse windows shrink.

openbao · DigiCert CertCentral

DETECT

Tie a packet to the process that sent it, without a separate host agent

Wireshark can record process attribution at capture time and expose it as a filter field, which collapses the usual two-tool dance of correlating a suspicious flow against EDR or netstat output after the fact. NETRESEC's FlowCarp adds a statistical classifier for protocol and malware family over a PCAP, with a hosted demo for a quick first pass.

Wireshark · NETRESEC

RESPOND

Run one triage queue across a fleet of serverless projects

Cross-Project Search lets an origin project pull alerts, ES|QL hunts, ML results and entity-store data from up to 100 linked projects running roughly 2,100 detection rules. Analysts stop tab-hopping between tenant projects to answer whether an indicator appeared anywhere else in the estate.

Elastic Security Labs

PROTECT

Box in a coding agent before it touches the repo

GitHub's Copilot app adds per-project local sandboxing that constrains file, network and credential access during agent sessions; Corridor puts a security scan at the pre-commit step for Devin, Codex, Claude Code and Cursor, so agent-written code is checked before it becomes history.

> Docker Sandboxes' `sbx settings` CLI makes the equivalent controls — kit signature checks, SSH agent forwarding, clipboard access — configurable rather than defaults you inherit.
>
> GitHub

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Wireless & RF

airgeddon

Feedback ⓘ v12.02 NOTES ↗ RANK

This is a multi-use bash script for Linux systems to audit wireless networks.

> airgeddon v12.02 extends DoS pursuit to 5GHz, adds nftables NAT compatibility, replaces its ISC DHCP server with Kea, and improves terminal window sizing.

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

01 **5GHz band hopping in DoS pursuit mode** NEW 45

DoS pursuit mode now supports 5GHz band hopping, enabling the tool to pursue targets that move across both 2.4GHz and 5GHz bands rather than only 2.4GHz.

– Explains behaviour change but no config/flag names given v12.02

02 **nftables NAT support validation** NEW 45

Adds iptables/nftables NAT support validation, extending compatibility to systems using the modern nftables firewall framework alongside existing iptables setups.

– Names the framework but no exact commands or checks v12.02

03 **Dynamic xterm sizing and positioning** IMPROVED 30

xterm windows are now sized and positioned dynamically based on actual desktop and window metrics, rather than fixed dimensions.

– Describes behaviour but no configurable options named v12.02

04 **Kea DHCP server replaces ISC DHCP** BREAKING 25

airgeddon v12.02 replaces the ISC DHCP server with Kea DHCP server for its DHCP functionality.

– Mentioned only in summary, no migration or config detail v12.02

v12.02 NOTES ↗ LATEST STABLE

airgeddon v12.02 adds 5GHz DoS pursuit, nftables NAT support, and replaces ISC DHCP with Kea DHCP server.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v12.02
https://github.com/v1s1t0r1sh3r3/airgeddon.git
# already have the repo? check out this version:
$ git checkout v12.02
```

Adds 5GHz band hopping support to DoS pursuit mode, enabling pursuit of targets across both 2.4GHz and 5GHz bands.

Adds iptables/nftables NAT support validation, extending compatibility to systems using the modern nftables firewall framework.

Adds dynamic xterm sizing and positioning based on actual desktop and window metrics.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

Invicti

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

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

> Invicti expanded API-driven access to discovery data, adding REST API retrieval for Website Discovery, spec source visibility in API Discovery, and richer filtering in the Website Discovery list.

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

01 **Filtering and interaction improvements in Website Discovery list**

IMPROVED 60

The Website Discovery list gains discovery date filtering, exact-match and exclusion text filters, and multi-select risk score filtering. Clicking a count in the list now filters results to show only entries matching that value.

— Lists all filter types and the click-to-filter behavior concretely.

changelog-20260924-d4de806d

THINNER COVERAGE BELOW

02 **Spec source location shown in API Discovery** NEW 50

API Discovery now displays the source location of each API specification, showing where each specification was found and which endpoints were auto-detected during a scan.

— Describes what changed but only a UI display, no exact navigation given.

changelog-20260924-d4de806d

03 **Website Discovery exposed via REST API** NEW 45

Website Discovery is now exposed through the Invicti Platform REST API, enabling retrieval of discovered websites and integration of discovery results into external workflows and cyber tooling.

— Names REST API surface but no endpoint path given. changelog-20260924-d4de806d

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Invicti adds REST API access for Website Discovery and surfaces API spec source locations in API Discovery.

Exposes Website Discovery via the Invicti Platform REST API, enabling retrieval of discovered websites and integration of discovery results into external workflows and cyber tooling.

Adds API specification source location display in API Discovery, showing where each specification was found and which endpoints were auto-detected during a scan.

Adds discovery date filtering, exact-match and exclusion text filters, and multi-select risk score filtering to the Website Discovery list.

Clicking a count in the Website Discovery list now filters the list to show only entries matching that value.

WAS THIS USEFUL?  

◆ SAST / SCA / Code Security

Aikido Security

Feedback ⓘ

🚀Code Coverage: find what your tests miss NOTES ↗ RANK

Aikido Security is a developer-first security platform that identifies and fixes vulnerabilities in code, dependencies, and infrastructure.

> Aikido shipped two major additions this window: a per-repository test coverage view integrated into code quality findings, and Deep Review, an agentic pull request reviewer that reasons across the full codebase.

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

01 **Code Coverage in Code Quality dashboard** NEW 82

Adds a `Code Quality > Code Coverage` section showing per-repository, per-file, and per-line test coverage alongside existing code quality findings. Accepts `LCOV` and `Cobertura` report formats with native support for Jest and pytest-cov outputs, and provides a GitHub Action and a GitLab CI component to upload reports from CI pipelines with no extra token or self-hosted component required. Supports merging split coverage reports from multiple services into a single view when multiple files are passed in one upload, integrates with GitHub, GitLab, Bitbucket, and Azure DevOps (plus a custom integration path for others), offers `Code Tree` and `File List` browsing modes, and tracks coverage over time with a per-repository trend graph updated on each upload.

Navigate to the per-file line coverage view in Aikido to check what tests already cover before editing a flagged file.

> 📍 1. In Aikido, go to Code Quality › Code Coverage. 2. Select the repository from the repo list. 3. Open the `File List` tab to sort files by coverage percentage. 4. Click the file of interest to see line-by-line coverage highlighting.

— Names formats, CI integrations and UI tabs; no exact CLI/YAML snippet given

🚀Code Coverage: find what your tests m…

02 **Deep Review agentic pull request reviewer** NEW 79

Deep Review is an agentic PR reviewer that reasons across the full codebase — including middleware, global config, and linked repositories, not just the diff — flagging bugs and unintended side effects in addition to security findings. It saves accepted-risk and false-positive decisions to memory so future reviews reflect prior feedback, and runs after existing PR checks like SAST and dependency checks to build on rather than repeat their findings. It's enabled per repository with configurable linked repos and a review level setting, billed at 1 credit ($1) per PR with unlimited re-reviews on the same PR and a configurable monthly spend cap set via Aikido Wallet.

Enable Deep Review on a repository and link related repos so cross-repo logic tracing covers your full stack before merging a sensitive change.

> 📍 In Aikido, navigate to your repository settings, enable Deep Review for the target repo, add linked repositories under `Related repos`, and select the review level that fits your team.

Cap monthly Deep Review spend to keep costs predictable across active repositories.

> 📍 In Aikido, go to Aikido Wallet › Deep Review and set a monthly credit cap. Each PR costs 1 credit ($1) with unlimited re-reviews on the same PR.

— Describes mechanism and pricing; navigation given but no API/config key

🚀Deep Review: the agentic PR review th…

🚀Code Coverage: find what your tests miss **NOTES ↗** **STABLE**

Aikido adds Code Coverage: line-level test coverage for every repo, paired with code quality findings, accepting LCOV and Cobertura reports.

└──▷ **HOW TO FIND IT**

Navigate to the per-file line coverage view in Aikido to check what tests already cover before editing a flagged file.

> 📍 1. In Aikido, go to Code Quality › Code Coverage. 2. Select the repository from the repo list. 3. Open the `File List` tab to sort files by coverage percentage. 4. Click the file of interest to see line-by-line coverage highlighting.

Adds `Code Quality > Code Coverage` section to the Aikido dashboard, showing per-repository, per-file, and per-line test coverage alongside existing code quality findings.

Accepts LCOV and Cobertura coverage report formats, with native support for outputs from Jest and pytest-cov.

Provides a GitHub Action and a GitLab CI component for uploading coverage reports to Aikido from CI pipelines, requiring no extra token or self-hosted component.

Supports merging split coverage reports across multiple services into a single view when multiple report files are passed in a single upload.

Integrates natively with GitHub, GitLab, Bitbucket, and Azure DevOps SCM platforms, with a custom integration path for other platforms.

+2 more

🚀Deep Review: the agentic PR review that cuts merge time and bugs shipped

NOTES ↗ LATEST STABLE

Aikido launches Deep Review, an agentic PR reviewer that traces logic across repos and posts findings with suggested fixes directly on pull requests.

└──▷ HOW TO FIND IT

Enable Deep Review on a repository and link related repos so cross-repo logic tracing covers your full stack before merging a sensitive change.

> 📍 In Aikido, navigate to your repository settings, enable Deep Review for the target repo, add linked repositories under `Related repos`, and select the review level that fits your team.

Cap monthly Deep Review spend to keep costs predictable across active repositories.

> 📍 In Aikido, go to Aikido Wallet › Deep Review and set a monthly credit cap. Each PR costs 1 credit ($1) with unlimited re-reviews on the same PR.

Introduces Deep Review, an agentic pull request reviewer that reasons across your full codebase — including middleware, global config, and linked repositories — not just the diff.

Flags bugs and unintended side effects in addition to security findings on every pull request.

Saves accepted-risk and false-positive decisions to memory so future reviews reflect prior feedback.

Runs after existing PR checks (SAST, dependency checks) to build on their findings rather than repeat them.

Enabled per repository with configurable linked repos and a review level setting; billed at 1 credit ($1) per PR with unlimited re-reviews on the same PR and a configurable monthly spend cap.

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ announced · Sep 24, 2026 BLOG ↗ RANK

Get started with Semgrep to help you catch, flag, and fix real vulnerabilities before they ship.

> Semgrep Supply Chain launches zero-day malware detection with automated, org-wide response policies, proactive blast-radius notifications, and a private-beta Malware Firewall that blocks malicious packages before they reach developer machines.

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

01 **Malware Firewall (private beta)** NEW 69

A Malware Firewall, embedded in Semgrep Guardian via Agent admin, intercepts dependency requests at the network level between the package manager and public registries and blocks malicious packages before they reach developer machines, requiring no per-device install. It is currently in private beta.

– Names admin location and mechanism though still private beta launch-20260924-77b2b3de

THINNER COVERAGE BELOW

02 **Early-notification opt-in for incident blast radius** NEW 59

An opt-in early-notification feature proactively informs customers via Slack or webhooks of their blast radius — affected repos, projects, findings, and scan progress — when a new supply chain incident is detected.

– Clear opt-in mechanism and named blast-radius fields launch-20260924-77b2b3de

03 **Automated Supply Chain response policies** NEW 58

Org-wide Supply Chain policies can trigger automated responses — opening a Jira ticket, sending a Slack alert, or calling a webhook — the moment a malware finding is generated, with no human in the loop.

– Names response channels but no config surface or setup steps launch-20260924-77b2b3de

04 **Zero-day supply chain incident detection** NEW 56

Semgrep now detects zero-day incidents by rescanning all customer repositories with AI-assisted, expert-verified advisory rules within minutes of a supply chain incident, posting findings to the Semgrep dashboard organized by advisories rolled up into incidents.

– Describes mechanism and timing but only points to a dashboard launch-20260924-77b2b3de

└──▷ **ALSO FROM THESE RELEASES**

announced · Sep 24, 2026 BLOG ↗

Semgrep Supply Chain adds zero-day malware detection with automated org-wide response policies and a malware firewall (private beta).

Adds org-wide Supply Chain policies that trigger automated responses — open a Jira ticket, send a Slack alert, or call a webhook — the moment a malware finding is generated, with no human in the loop.

Adds zero-day incident detection that rescans all customer repositories with AI-assisted, expert-verified advisory rules within minutes of a supply chain incident and posts findings to the Semgrep dashboard organized by advisories rolled up into incidents.

Adds early-notification opt-in that proactively informs customers via Slack or webhooks of their blast radius — affected repos, projects, findings, and scan progress — when a new supply chain incident is detected.

Adds a Malware Firewall (private beta), embedded in Semgrep Guardian via Agent admin, that intercepts dependency requests at the network level between the package manager and public registries and blocks malicious packages before they reach developer machines, with no per-device install.

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ

More ways to request and configure Copilot code reviews (2026-09-23) 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 finer control over Copilot code reviews — with personal and enterprise-wide default effort levels — and added per-project local sandboxing to the Copilot app to restrict file, network, and credential access during agent sessions.

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

01 Local sandboxing for Copilot app sessions NEW

93

The GitHub Copilot app adds per-project sandbox settings covering filesystem access (additional read/write, read-only, and denied folder lists), outbound internet and local network controls, and Git/GitHub CLI credential access for local repository and working tree sessions. A new `/sandbox on` command enables sandboxing mid-session without changing the project default, and if the OS cannot enforce the requested policy the sandboxed shell errors out rather than silently running unsandboxed.

Enable sandboxing mid-session on an active local Copilot session without altering the project-level default.

```
$ /sandbox on
```

Turn on sandboxing for all new sessions in a project via the app settings UI.

📍 1. Open the GitHub Copilot app settings. 2. Select your project. 3. Under Sandbox, enable Sandbox new sessions. 4. Configure filesystem (read/write, read-only, denied folders), network (outbound internet, local network), and credential (Git credentials, GitHub CLI credentials) policies as needed.

— Names runnable command and exact policy categories

Local sandboxing in the GitHub Copilot …

02 Configurable Copilot code review effort levels NEW

80

A new `code review` page under profile → Copilot settings (available on every plan, including Business and Enterprise) adds separate toggles for automatic reviews on PR creation, new pushes, and draft PRs you create or coauthor, plus a personal `default review effort` setting (`Lite` or `Balanced`) that applies to all requested and automatic reviews while still allowing per-request overrides from the `Reviewers` panel. Enterprise administrators can also set an enterprise-wide default review effort (`Lite` , `Balanced` , or GitHub default) for organization-owned repositories, with organization- and repository-level overrides still permitted.

Enable automatic Copilot reviews — including on new pushes and draft PRs — and set a default effort level so every PR you open gets reviewed at your chosen depth without manual configuration.

> 📍 In GitHub, go to your profile → Copilot settings → Code review. Enable `Automatic reviews` , toggle on `New pushes` and `Draft pull requests` , then set `Default review effort` to Balanced.

Standardize Copilot review depth across all org-owned repositories in an enterprise, so teams without their own override inherit a consistent baseline.

> 📍 As an enterprise administrator, go to your enterprise settings → Copilot → Code review, and set the enterprise-wide `Default review effort` to Lite, Balanced, or GitHub default.

— Exact settings page and toggle names, no API surface

More ways to request and configure Copi…

More ways to request and configure Copilot code reviews (2026-09-23) BLOG ↗

LATEST STABLE

GitHub Copilot code review adds per-user effort controls on all plans and an enterprise-wide default effort setting for org repos.

└──▷ **HOW TO FIND IT**

Enable automatic Copilot reviews — including on new pushes and draft PRs — and set a default effort level so every PR you open gets reviewed at your chosen depth without manual configuration.

> 📍 In GitHub, go to your profile → Copilot settings → Code review. Enable `Automatic reviews` , toggle on `New pushes` and `Draft pull requests` , then set `Default review effort` to Balanced.

Standardize Copilot review depth across all org-owned repositories in an enterprise, so teams without their own override inherit a consistent baseline.

📍 As an enterprise administrator, go to your enterprise settings → Copilot → Code review, and set the enterprise-wide `Default review effort` to Lite, Balanced, or GitHub default.

Adds a dedicated `code review` page under profile → Copilot settings, available on every Copilot plan (including Business and Enterprise), with separate toggles for automatic reviews on PR creation, new pushes, and draft PRs you create or coauthor.

Adds a personal `default review effort` setting (`Lite` or `Balanced`) on the Copilot code review settings page, applying to all reviews you request — including automatic ones — while still allowing per-request overrides from the `Reviewers` panel.

Adds an enterprise-wide default review effort setting (`Lite` , `Balanced` , or GitHub default) for organization-owned repositories, configurable by enterprise administrators, with organization- and repository-level overrides still permitted.

`Local sandboxing in the GitHub Copilot app (2026-09-23)` **BLOG ↗** **STABLE**

GitHub Copilot app gains per-project local sandboxing to restrict file, network, and credential access during agent sessions.

└──▷ **TRY IT**

Enable sandboxing mid-session on an active local Copilot session without altering the project-level default.

```
$ /sandbox on
```

Turn on sandboxing for all new sessions in a project via the app settings UI.

📍 1. Open the GitHub Copilot app settings. 2. Select your project. 3. Under `Sandbox` , enable `Sandbox new sessions` . 4. Configure filesystem (read/write, read-only, denied folders), network (outbound internet, local network), and credential (Git credentials, GitHub CLI credentials) policies as needed.

Adds `/sandbox on` command to enable sandboxing for an already-running local session without changing the project default.

Adds per-project sandbox settings in the GitHub Copilot app covering filesystem (additional read/write, read-only, and denied folder lists), outbound internet and local network controls, and Git/GitHub CLI credential access for local repository and working tree sessions.

Enforces hard failure — the sandboxed shell errors out rather than silently running without a sandbox — when the OS cannot enforce the requested policy.

WAS THIS USEFUL? 👍 👎

Snyk Code

Feedback ⓘ changelog · Sep 23, 2026 BLOG ↗ RANK

Snyk Code is a static analysis security tool that scans source code to identify and fix vulnerabilities, bugs, and code quality issues.

> Snyk Code added a new command that wires Snyk scanning directly into AI coding tools, checking generated code for vulnerabilities in the background as it's written.

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

01 **Snyk Studio install for AI coding tools** NEW 80

The new `snyk studio install --experimental` command configures Snyk Studio inside AI coding tools including Cursor, Claude Code, Codex, Copilot, Gemini, Kiro, and Windsurf. Once installed, it scans code generated by these tools in the background as it is written.

Set up Snyk Studio in your AI coding environment so code generated by tools like Cursor or Copilot is scanned for vulnerabilities in the background as you work.

```
$ snyk studio install --experimental
```

— Names exact command and supported tools, but marked experimental. 2026-09-23-5abc9ed0

changelog · Sep 23, 2026 BLOG ↗ LATEST STABLE

New `snyk studio install --experimental` command integrates Snyk scanning directly into AI coding tools like Cursor, Copilot, and Claude Code.

└──▷ **TRY IT**

Set up Snyk Studio in your AI coding environment so code generated by tools like Cursor or Copilot is scanned for vulnerabilities in the background as you work.

```
$ snyk studio install --experimental
```

WAS THIS USEFUL? 👍 👎

Snyk Open Source

Feedback ⓘ changelog · Sep 23, 2026

BLOG ↗ RANK

Snyk Open Source scans and fixes security vulnerabilities in open source dependencies across development projects.

> Snyk Open Source added filtering controls to `snyk fix --agentic` and changed its default failure-handling behavior to preserve partial changes instead of auto-reverting them.

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

01 **Filters for agentic Open Source fixes** NEW 78

`snyk fix --agentic` gains three new flags: `--severity-filter` to limit automated fixes to chosen severity levels, `--breakability-filter` to restrict Open Source upgrades to those matching a chosen breakability level, and `--exclude-ids` to skip specific issue IDs while fixing everything else.

Fix only critical and high severity issues in CI, skipping known-noisy issue IDs, so the agentic run stays focused on what matters most.

```
$ snyk fix --agentic --severity-filter=critical,high --exclude-ids=SNYK-JS-12345,SNYK-JS-67890
```

Restrict agentic Open Source upgrades to low-breakability changes only, reducing the chance of breaking builds in a sensitive production repo.

```
$ snyk fix --agentic --breakability-filter=low
```

— Three named flags with runnable commands, but no mechanism detail beyond the flag purpose

2026-09-23-f0197e10

02 **Agentic fix revert behavior changed to opt-in** BREAKING 75

`snyk fix --agentic` no longer automatically reverts failed fix attempts on upgrade; changes are now kept in place for review by default. Add `--enable-revert` to restore the previous automatic rollback behavior.

Investigate why an agentic fix failed without losing the partial changes — the new default keeps them in place; add `--enable-revert` only when you want the old auto-rollback behavior back.

```
$ snyk fix --agentic --enable-revert
```

— Clear before/after behavior and exact flag to restore old default `2026-09-23-f0197e10`

└──▷ **BREAKING ON UPGRADE**

! `snyk fix --agentic` no longer automatically reverts failed fix attempts on upgrade — changes are now kept for review by default. Add `--enable-revert` to restore the previous automatic rollback behavior.

changelog · Sep 23, 2026 | BLOG ↗ | LATEST | STABLE

snyk fix `--agentic` gains severity, breakability, and exclusion filters plus opt-in revert control

└──▷ **TRY IT**

Fix only critical and high severity issues in CI, skipping known-noisy issue IDs, so the agentic run stays focused on what matters most.

```
$ snyk fix --agentic --severity-filter=critical,high --exclude-ids=SNYK-JS-12345,SNYK-JS-67890
```

Restrict agentic Open Source upgrades to low-breakability changes only, reducing the chance of breaking builds in a sensitive production repo.

```
$ snyk fix --agentic --breakability-filter=low
```

Investigate why an agentic fix failed without losing the partial changes — the new default keeps them in place; add `--enable-revert` only when you want the old auto-rollback behavior back.

```
$ snyk fix --agentic --enable-revert
```

Adds `--severity-filter` flag to `snyk fix --agentic` to limit automated fixes to only the chosen severity levels.

Adds `--breakability-filter` flag to `snyk fix --agentic` to restrict Open Source upgrades to those matching the chosen breakability level.

Adds `--exclude-ids` flag to `snyk fix --agentic` to skip specific issue IDs while fixing everything else.

└──▷ **BREAKING ON UPGRADE**

! `snyk fix --agentic` no longer automatically reverts failed fix attempts on upgrade — changes are now kept for review by default. Add `--enable-revert` to restore the previous automatic rollback behavior.

WAS THIS USEFUL?  

Black Duck

Feedback ⓘ

Polaris release update: Bringing AI security findings, scan governance, and remediation

BLOG ↗ RANK

Black Duck helps organizations secure their software supply chain by providing deep visibility into open source components, licenses, and vulnerabilities.

> Polaris adds AI-generated Signal AI findings with bring-your-own-LLM configuration, organization-wide rule profiles, expanded on-prem SCM support, async SCA vulnerability alerts, and a dynamic scan sitemap view.

└──▷ **WHAT SHIPPED · 7 FEATURES** most completely described first ? what's the number?

01 **BYO LLM configuration for Signal AI agents** NEW 75

Supports per-agent LLM configuration across each of Signal AI's six agents, with provider, model, and endpoint independently selectable — supporting OpenAI, Azure OpenAI, AWS Bedrock, GCP Vertex AI, or any OpenAI-compatible proxy with automatic provider detection.

— Names all supported providers and per-agent scope

Polaris release update: Bringing AI sec…

02 **Central rule configuration with hierarchical inheritance** NEW 65

Introduces central rule configuration in Polaris, letting security teams define rule profiles once and apply them at the organization, application, or project level via hierarchical inheritance, automatically enforced across CI pipelines, CLI scans, and integrations.

— Explains mechanism and enforcement scope, no exact config key

Polaris release update: Bringing AI sec…

03 **On-prem SCM support extended to GitLab, Azure DevOps Server, Bitbucket**

NEW

65

Extends on-premises SCM support to self-hosted GitLab, Azure DevOps Server, and Bitbucket Server/Data Center (in addition to previously released self-managed GitHub), enabling repository discovery, bulk onboarding, SAST/SCA scan triggers, and pull request–driven remediation workflows.

— Names each SCM platform and the workflows enabled

Polaris release update: Bringing AI sec…

04 **Async SCA email alerts for critical/high vulnerabilities** NEW

60

Adds async SCA notifications — email alerts for new critical- or high-severity open source vulnerabilities, or reclassifications of existing vulnerabilities to critical or high, with project- and branch-level subscription controls.

— Names trigger conditions and subscription scope

Polaris release update: Bringing AI sec…

THINNER COVERAGE BELOW

05 **Signal AI findings unified with deterministic scans** NEW

55

Polaris now accepts SARIF-format upload of Signal AI code analysis results, classifying them under a new AI tool type and managing them alongside deterministic findings through unified policies, triage, audit trails, and governance workflows.

— Names SARIF format and workflow scope but no exact steps

Polaris release update: Bringing AI sec…

06 **2025 CWE Top 25 added to policy criteria** NEW

50

Adds support for the 2025 CWE Top 25 (alongside the existing 2024 list) across policy filters and violation criteria, dashboards, reporting, and issue management.

— Names the list and surfaces it touches, no config detail

Polaris release update: Bringing AI sec…

07 **fAST Dynamic sitemap view** NEW

45

Adds a fAST Dynamic sitemap view that displays discovered URLs, endpoints, and resources in a hierarchical map so teams can validate dynamic testing coverage and identify gaps or unexpected endpoints.

— Describes purpose but no exact navigation path given

Polaris release update: Bringing AI sec…

Polaris release update: Bringing AI security findings, scan governance, and remediatio

BLOG ↗ LATEST STABLE

Polaris gains BYO LLM for Signal AI analysis, central scan rule profiles, on-prem SCM support, async SCA alerts, and a dynamic sitemap view.

Adds SARIF-format upload of Signal AI code analysis results into Polaris, where they are classified under the new AI tool type and managed alongside deterministic findings through unified policies, triage, audit trails, and governance workflows.

Supports per-agent LLM configuration across each of Signal AI's six agents, with provider, model, and endpoint independently selectable — supporting OpenAI, Azure OpenAI, AWS Bedrock, GCP Vertex AI, or any OpenAI-compatible proxy with automatic provider detection.

Introduces central rule configuration in Polaris, letting security teams define rule profiles once and apply them at the organization, application, or project level via hierarchical inheritance, automatically enforced across CI pipelines, CLI scans, and integrations.

Adds support for the 2025 CWE Top 25 (alongside the existing 2024 list) across policy filters and violation criteria, dashboards, reporting, and issue management.

Extends on-premises SCM support to self-hosted GitLab, Azure DevOps Server, and Bitbucket Server/Data Center (in addition to previously released self-managed GitHub), enabling repository discovery, bulk onboarding, SAST/SCA scan triggers, and pull request–driven remediation workflows.

+2 more

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Veracode is a cloud-based application security platform that performs static, dynamic, and software composition analysis to identify vulnerabilities in code.

> Veracode's biggest changes this window are a sweeping expansion of supported languages, frameworks and platform versions across its Static Analysis Engine, alongside new `veracode.json` exclusion controls, a 200 MB Pipeline Scan artifact limit increase, and a redesigned EASM interface.

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

01 **Scan exclusions via veracode.json** NEW 70

Static Analysis now supports excluding Java `.jar` files and excluding specific Java classes from scans through the `veracode.json` configuration file.

— Names the exact config file readers can edit today.

Static Analysis updates (2026-09-23)

02 **Python scanner analysis engine enhancements** IMPROVED 65

The Python scanner gains improved taint tracking, cross-file analysis, and pattern recognition, plus added support for analyzing PySpark, Flask, Django, and FastAPI applications.

— Names mechanisms and frameworks but no direct user action.

Static Analysis updates (2026-09-23)

03 **Taint analysis coverage expansion across Python, Swift and .NET**

IMPROVED 60

Adds taint sinks for the `send_from_directory` and `send_file` functions in Python, adds taint propagation support for the Swift standard library, and expands `.NET` taint sink coverage for CWE-80 to detect `System.Web.UI.Page.Title` and `System.Web.UI.HtmlControls.HtmlHead.Title`.

— Names exact sinks and APIs but is an internal detection improvement.

Static Analysis updates (2026-09-23)

THINNER COVERAGE BELOW

04 **Expanded language, framework and platform support** NEW 55

Static Analysis added support for iOS 27, Go 1.27, Android 17, Dart 3.10/Flutter 3.38, Dart 3.11/Flutter 3.41, Dart 3.12/Flutter 3.44, Dart 3.13/Flutter 3.47, Kotlin 2.3 and 2.4, JDK 25 (enhanced), Xcode 26.4, React 19, Angular 20, Next.js 15.x and 16.x, the Svelte framework, Ruby 4.0 and Rails 8.1, Apex 64 and 65, Spring Boot 4/Spring Core 7/Spring Security 7, the Go Fiber web framework, Fast HTTP in Go, Psycopg 3 in Python, PyYAML, AWS Glue in Python, the Python `google-cloud-bigquery` library, and Apple CloudKit and Security framework APIs.

— Long list of named versions but no mechanism beyond 'now supported'.

Static Analysis updates (2026-09-23) changelog-20260924-9c136d32

---

05 **Pipeline Scan support for artifacts over 200 MB** IMPROVED

55

Pipeline Scan now supports scanning packaged artifacts larger than 200 MB, with no changes to existing integrations required.

— States the new size limit and that integrations need no changes.

Static Analysis updates (2026-09-23)

---

06 **Static Analysis Engine v2026.09 detection accuracy improvements**

IMPROVED

53

The Static Analysis Engine v2026.09 improves detection of CWE-259 and CWE-798 flaws, CWE-89 flaws in PySpark applications (with improved SQLAlchemy analysis), and CWE-200 flaws, all reducing false positives, and improves analysis of `CALL` commands to enhance CWE-252 detection.

— Names specific CWEs and mechanism but no user-facing action.

changelog-20260924-9c136d32

---

07 **Hardcoded secrets detection for environment config files**

IMPROVED

50

Hardcoded secrets detection was expanded to include application configuration files with `-local` and `-prod` naming patterns.

— Names the exact file naming patterns now scanned. Static Analysis updates (2026-09-23)

---

08 **Java source code scanning and viewing** NEW

30

Adds support for scanning Java source code and for viewing Java source code directly on the platform.

— Thin description with no further mechanism given. Static Analysis updates (2026-09-23)

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Veracode: Adds support for iOS 27 in the Static Analysis Engine v2026.09.

Adds support for Go 1.27 in the Static Analysis Engine v2026.09.

Redesigns the EASM interface with improved navigation and usability, with a legacy UI fallback option.

Improves detection of CWE-259 and CWE-798 flaws in the Static Analysis Engine v2026.09, reducing false positives.

Improves detection of CWE-89 flaws in PySpark applications and improves SQLAlchemy analysis to reduce CWE-89 false positives in the Static Analysis Engine v2026.09.

Improves detection of CWE-200 flaws in the Static Analysis Engine v2026.09, reducing false positives.

+1 more

Static Analysis updates (2026-09-23) BLOG ↗ STABLE

Veracode Static Analysis adds iOS 27, Go 1.27, DryRun Security integration, and Pipeline Scan support for artifacts larger than 200 MB

Adds support for excluding Java `.jar` files from static analysis using `veracode.json`.

Adds support for excluding specific Java classes from Static Analysis through the `veracode.json` configuration file.

Pipeline Scan now supports scanning packaged artifacts larger than 200 MB with no changes to existing integrations required.

Adds support for Dart 3.13 and Flutter 3.47.

Adds support for Android 17.

+30 more

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

Truffle Security trufflehog

Feedback ⓘ v3.97.9

NOTES ↗ CODE ↗ RANK

trufflehog discovers, classifies, and validates leaked credentials

> TruffleHog's latest releases add Okta OAuth 2.0 client credential detection, per-detector verification latency metrics, and more resilient handling of verification failures across 45 detectors, alongside smaller job-progress and SharePoint metadata additions.

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

01 **Okta OAuth 2.0 client credential detection** NEW 78

The Okta detector now detects and verifies OAuth 2.0 client credentials (`client_id` + `client_secret`) in addition to the existing SSWS API token detection. Verification uses the `client_credentials` grant against `https://{domain}/oauth2/default/v1/token` .

– Names exact grant type and verification endpoint. v3.97.8

02 **Per-detector verification latency metrics** NEW 62

Adds per-detector verification wall-time tracking to the verification cache via an optional `DetectorMetricsReporter` interface, enabling metrics backends to record remote verification latency broken down by detector type.

– Names interface but no usage example given. v3.97.8

THINNER COVERAGE BELOW

03 **Findings preserved despite verification errors** IMPROVED 55

45 secret detectors now report findings even when live verification fails for operational reasons (request build errors, transport errors, body read errors) rather than silently discarding them, using `SetVerificationError` to record the failure.

– Names mechanism and count but no user-facing action. v3.97.8

04 **Job progress error inspection method** NEW 55

Adds an Err() method to `JobProgress` and `JobProgressRef` , mirroring `context.Err` , so callers can programmatically read why a scan job ended.

– Names exact API but no example of use. v3.97.9

05 **SharePoint metadata version and comment fields** IMPROVED

Adds `version` and `comment_number` fields to SharePoint source metadata.

— Bare field addition with no further detail. v3.97.9

v3.97.9 NOTES ↗ CODE ↗ LATEST STABLE

TruffleHog v3.97.9 adds version and comment_number fields to SharePoint source metadata.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.97.9
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.97.9
```

Adds Err() method to `JobProgress` and `JobProgressRef`, mirroring `context.Err` so callers can programmatically read why a scan job ended.

v3.97.8 NOTES ↗ CODE ↗ STABLE

TruffleHog v3.97.8 adds Okta OAuth 2.0 client credential detection and per-detector verification timing metrics.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.97.8
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.97.8
```

Adds detection and verification for Okta OAuth 2.0 client credentials (`client_id` + `client_secret`) in the Okta detector, complementing the existing SSWS API token detection; verifies via the `client_credentials` grant against `https://{domain}/oauth2/default/v1/token`.

Adds per-detector verification wall-time tracking to the verification cache via an optional `DetectorMetricsReporter` interface, enabling metrics backends to record remote verification latency broken down by detector type.

Ensures 45 secret detectors now report findings even when live verification fails for operational reasons (request build errors, transport errors, body read errors) rather than silently discarding them, using `SetVerificationError` to record the failure.

WAS THIS USEFUL?  

openbao

Feedback ⓘ v2.7.0 NOTES ↗ RANK

OpenBao is a software solution to manage, store, and distribute sensitive data including secrets, certificates, and keys.

> OpenBao's v2.7.0 release is dominated by post-quantum readiness (PQC TLS key exchange and ML-DSA signing), External Keys support for KMS-backed PKI and Transit operations, PostgreSQL read scaling, a new PebbleDB storage backend, and Control Groups, alongside removal of several built-in seals, the Kerberos/LDAP auth engines, and the file storage backend as external-plugin-only; v2.6.3 and the Proxy docs add smaller config, SDK, and environment-variable surfaces.

└──▷ **WHAT SHIPPED · 26 FEATURES** most completely described first ? what's the number?

01 **Post-quantum TLS key exchange enforcement** NEW 90

Adds `tls_key_exchange_preferences` to server, agent, and proxy listener config to enforce pure post-quantum TLS key exchange algorithms (`SecP256r1MLKEM768` , `SecP384r1MLKEM1024` , `X25519MLKEM768` , `MLKEM1024`).

Enforce post-quantum key exchange on a listener so all inbound TLS connections use PQC algorithms.

hcl

```
listener "tcp" {
  address = "0.0.0.0:8200"
  tls_cert_file = "/etc/openbao/tls/server.crt"
  tls_key_file  = "/etc/openbao/tls/server.key"
  tls_key_exchange_preferences = "X25519MLKEM768,MLKEM1024"
}
```

— Config key, algorithm names, and a runnable listener config example. v2.7.0

02 **Read consistency controls via listener config and request headers**

NEW 90

Adds `consistency_fallback_behavior` , `consistency_missing_header_forward` , and `consistency_max_index_wait` listener configuration parameters, and supports `X-Vault-Index` and `X-Vault-Inconsistent` request headers (values `fail` , `forward-active-node` , `await-state`) to enforce client/server read consistency.

Hold a read request on a standby node until it catches up to the index returned by a prior write, falling back to 429 if the wait expires — ensuring read-your-writes consistency.

```
$ curl -s -H "X-Vault-Token: $BAO_TOKEN" \
  -H "X-Vault-Index: $WRITE_INDEX" \
  -H "X-Vault-Inconsistent: await-state" \
  https://openbao.example.com:8200/v1/secret/data/myapp
```

— Named config keys, header values, and a runnable curl example. v2.7.0

03 **Environment variable configuration for OpenBao Proxy** NEW

90

Adds `BAO_PROXY_DISABLE_IDLE_CONNECTIONS` and `BAO_PROXY_DISABLE_KEEP_ALIVES` to control idle connection and keep-alive behaviour, `BAO_ADDR` to specify the upstream BAO server address, `BAO_CACERT` / `BAO_CAPATH` for CA certificate verification, `BAO_CLIENT_CERT` / `BAO_CLIENT_KEY` for mutual TLS, `BAO_SKIP_VERIFY` to disable TLS verification, `BAO_TLS_SERVER_NAME` to override the TLS server name, and `BAO_MAX_RETRIES` to configure maximum retries in OpenBao Proxy.

Run OpenBao Proxy with mutual TLS and a capped retry count — useful when fronting a hardened BAO cluster in a CI pipeline.

```
$ export BAO_ADDR=https://bao.internal:8200
export BAO_CACERT=/etc/bao/ca.pem
export BAO_CLIENT_CERT=/etc/bao/client.crt
export BAO_CLIENT_KEY=/etc/bao/client.key
export BAO_MAX_RETRIES=3
bao proxy -config=/etc/bao/proxy.hcl
```

Disable idle connections and keep-alives in OpenBao Proxy when operating behind a load balancer that enforces strict connection limits.

```
$ export BAO_PROXY_DISABLE_IDLE_CONNECTIONS=true
export BAO_PROXY_DISABLE_KEEP_ALIVES=true
bao proxy -config=/etc/bao/proxy.hcl
```

— Every env var named with two runnable shell examples. product docs

04 **External Keys support for PKI and Transit via KMS plugins**

NEW

85

PKI and Transit engines can now use KMS plugins via `/sys/external-keys` APIs to perform cryptographic operations without storing key material, with built-in Transit-backed key support and a `kms-pkcs11` plugin for PKCS#11. KMS plugin configuration stanzas are reloaded on `SIGHUP`, respawning plugin clients as needed to enable runtime plugin upgrades.

— Named endpoint and plugin mechanism, but no runnable example. v2.7.0

05 **OCI plugin management: version pinning, pruning, and auto-detection**

NEW

85

Adds `plugin_version=latest` to allow auth and secret mounts to automatically request the latest installed version of an external plugin, a `bao plugin prune` command to remove unused OCI plugin downloads from disk, support for pinning OCI-based plugins by manifest digest (e.g. `ghcr.io/openbao/openbao-plugin-secrets-aws@sha256:<digest>`) instead of a `sha256sum`, and automatic determination of `binary_name` for OCI plugins from the image's `ENTRYPOINT` or `CMD` when unset.

— Names a runnable command and exact config values. v2.7.0

06 **ML-DSA post-quantum signing in Transit and PKI engines** NEW

75

Introduces ML-DSA (`mldsa-44`, `mldsa-65`, `mldsa-87`) key types in the Transit engine for post-quantum signing and verification, and adds ML-DSA (NIST FIPS 204) support in the PKI engine for all CA, CSR, and leaf certificate operations.

— Named key types and scope but no command example. v2.7.0

07 **PostgreSQL horizontal read scalability** NEW

70

Adds PostgreSQL horizontal read scalability on the PostgreSQL storage backend, requiring `ha_enabled = true` and a reachable `cluster_addr`.

— Named config keys but no full example. v2.7.0

08 **Field and format output flags on bao CLI commands** IMPROVED

70

Adds a `-field` argument to `bao status`, `bao namespace seal-status` (which also gains `-format`), and `bao token lookup`.

— Exact flags and commands named, no worked example. v2.7.0

09 **Envoy XFCC client-cert decoder for auth/cert** NEW 65

Adds `x_forwarded_for_client_cert_decoders` Envoy decoder in `auth/cert` to process leaf certificates from the Envoy XFCC header.

— Named config field and target header, no example. v2.7.0

10 **PSS signature support for PKI intermediate CSRs** NEW 65

Supports `use_pss` in `issuers/generate/intermediate/:type` PKI API calls for CSRs with PSS signatures from RSA-typed intermediates.

— Named endpoint and parameter, no request example. v2.7.0

11 **Automatic TLS certificate reload** NEW 65

Adds `tls_auto_reload` configuration option to automatically reload TLS certificate and key files on change without requiring `SIGHUP`.

— Named config key and before/after behaviour. v2.7.0

12 **Built-in KMS-based seals moved to external plugins** BREAKING 60

The `pkcs11`, `alicloudkms`, `awskms`, `azurekeyvault`, `gcpckms`, and `ocikms` seals are no longer built-in and must be installed as external plugins.

— Names all affected seals but no migration steps. v2.7.0

13 **Kerberos and LDAP auth engines removed from core binary** 60

BREAKING

The built-in Kerberos auth engine (`auth/kerberos`) and built-in LDAP auth engine (`auth/ldap`) have been removed from the main OpenBao binary; both are now available in openbao-plugins.

— Names both engines and where they moved. v2.7.0

14 **X-Vault-Request header required for quit/cache-clear endpoints** 60

BREAKING

The `.../quit` and `.../cache-clear` agent/proxy endpoints now require the `X-Vault-Request` header when `require_request_header = true` is set in listener configuration; calls omitting the header are rejected.

— Named endpoints, header, and config condition. v2.7.0 v2.6.3

THINNER COVERAGE BELOW

15 **Control Groups ACL policy stanza** NEW — 55

Adds a `control_group` ACL policy stanza requiring second-party approval for requests to a path.

— Names the stanza but no config example. v2.7.0

16 **ACME client port and host configuration in PKI** NEW — 55

Adds optional ACME client port and host configuration in PKI to support non-privileged operation behind a load balancer on ports 80 and 443.

— Describes purpose but no named config field. v2.7.0

17 **Sanitized server config output gains new fields** IMPROVED — 50

Adds `disable_standby_reads`, `allow_unauthenticated_workflows`, and `unsafe_relative_paths` fields to the `command/server` sanitized config output.

— Named fields but only affects diagnostic output. v2.6.3

18 **ResolvePathOperation in SDK logical package** NEW — 50

Introduces `ResolvePathOperation` in `sdk/logical` to report canonical URLs for resources prior to ACL evaluation.

— Named SDK type and purpose, no usage example. v2.6.3

19 **PebbleDB storage backend** NEW — 45

Adds PebbleDB storage backend as a non-HA, durable backend using PebbleDB for transactions and improved performance.

— Names the backend but no config or benchmark detail. v2.7.0

20 **Reading backed-up shares via sys/raw** NEW — 45

Allows reading backed-up unseal or recovery shares via `sys/raw`.

— Named endpoint, minimal further detail. v2.7.0

21 **UI toggle to show internal engines** NEW — 45

Adds a Show internal engines button to the UI secrets list to hide the `cubbyhole` engine by default.

— Names UI location, no other mechanism. v2.7.0

22 **File storage backend removed** BREAKING — 45

The `file` storage backend has been removed from `command/server`.

— Names the backend but no migration guidance. v2.7.0

---

23 **plugin_auto_register now defaults to true** BREAKING

45

`plugin_auto_register` now defaults to `true`, which may cause previously unregistered plugins to be auto-registered on server start.

— Names the config key and its consequence. v2.7.0

---

24 **Auto Seal health check configuration** NEW

35

Adds configuration options to disable Auto Seal health checks and customize health check intervals and timeout.

— No config key names given. v2.7.0

---

25 **Sealed-state metric reporting** NEW

35

Reports a `vault.core.unsealed=0` metric via a dedicated metrics loop when OpenBao is sealed.

— Names the metric but no integration detail. v2.7.0

---

26 **revoke-self support for wrapping tokens** NEW

25

Adds support for `revoke-self` with wrapping tokens.

— Bare one-line addition with no mechanism. v2.7.0

└──▷ BREAKING ON UPGRADE

! The `pkcs11`, `alicloudkms`, `awskms`, `azurekeyvault`, `gcpckms`, and `ockms` seals are no longer built-in and must be installed as external plugins.

! The built-in Kerberos auth engine (`auth/kerberos`) has been removed from the main OpenBao binary; it is now in openbao-plugins.

! The built-in LDAP auth engine (`auth/ldap`) has been removed from the main OpenBao binary; it is now in openbao-plugins.

! The `file` storage backend has been removed from `command/server`.

! The `.../quit` and `.../cache-clear` agent/proxy endpoints now require the `X-Vault-Request` header when `require_request_header = true` is set.

! `plugin_auto_register` now defaults to `true`, which may cause previously unregistered plugins to be auto-registered on server start.

! The `agent` / `proxy` `.../quit` and `.../cache-clear` endpoints now require the `X-Vault-Request` request header when `require_request_header = true` is set in listener configuration — calls omitting the header will be rejected.

v2.7.0 NOTES ↗ LATEST STABLE

OpenBao v2.7.0 adds External KMS key support, ML-DSA post-quantum crypto, PQC TLS, PostgreSQL read scaling, Control Groups, and PebbleDB storage.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.7.0 https://github.com/openbao/openbao.git
# already have the repo? check out this version:
$ git checkout v2.7.0
```

└──▷ **USE IT**

Enforce post-quantum key exchange on a listener so all inbound TLS connections use PQC algorithms.

/etc/openbao/config.hcl

```
listener "tcp" {
  address = "0.0.0.0:8200"
  tls_cert_file = "/etc/openbao/tls/server.crt"
  tls_key_file  = "/etc/openbao/tls/server.key"
  tls_key_exchange_preferences = "X25519MLKEM768,MLKEM1024"
}
```

Hold a read request on a standby node until it catches up to the index returned by a prior write, falling back to 429 if the wait expires — ensuring read-your-writes consistency.

```
$ curl -s -H "X-Vault-Token: $BAO_TOKEN" \
  -H "X-Vault-Index: $WRITE_INDEX" \
  -H "X-Vault-Inconsistent: await-state" \
  https://openbao.example.com:8200/v1/secret/data/myapp
```

Adds `tls_key_exchange_preferences` to server, agent, and proxy listener config to enforce pure post-quantum TLS key exchange algorithms (`SecP256r1MLKEM768` , `SecP384r1MLKEM1024` , `X25519MLKEM768` , `MLKEM1024`).

Adds `consistency_fallback_behavior` , `consistency_missing_header_forward` , and `consistency_max_index_wait` listener configuration parameters for strong consistency control.

Supports `X-Vault-Index` and `X-Vault-Inconsistent` request headers (values: `fail` , `forward-active-node` , `await-state`) to enforce client/server read consistency.

Adds `plugin_version=latest` to allow auth and secret mounts to automatically request the latest installed version of an external plugin.

Adds `bao plugin prune` command to remove unused OCI plugin downloads from

disk.

+21 more

└──▷ **BREAKING ON UPGRADE**

- ! The `pkcs11`, `alicloudkms`, `awskms`, `azurekeyvault`, `gcpckms`, and `ocikms` seals are no longer built-in and must be installed as external plugins.
- ! The built-in Kerberos auth engine (`auth/kerberos`) has been removed from the main OpenBao binary; it is now in openbao-plugins.
- ! The built-in LDAP auth engine (`auth/ldap`) has been removed from the main OpenBao binary; it is now in openbao-plugins.
- ! The `file` storage backend has been removed from `command/server`.
- ! The `.../quit` and `.../cache-clear` agent/proxy endpoints now require the X-Vault-Request header when `require_request_header = true` is set.
- ! `plugin_auto_register` now defaults to `true`, which may cause previously unregistered plugins to be auto-registered on server start.

v2.6.3 NOTES ↗ STABLE

OpenBao v2.6.3 adds `ResolvePathOperation` to the SDK and exposes three new fields in sanitized server config output.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.6.3 https://github.com/openbao/openbao.git
# already have the repo? check out this version:
$ git checkout v2.6.3
```

Adds `disable_standby_reads`, `allow_unauthenticated_workflows`, and `unsafe_relative_paths` fields to the `command/server` sanitized config output.

Introduces `ResolvePathOperation` in `sdk/logical` to report canonical URLs for resources prior to ACL evaluation.

└──▷ **BREAKING ON UPGRADE**

- ! The `agent` / `proxy` `.../quit` and `.../cache-clear` endpoints now require the X-Vault-Request request header when `require_request_header = true` is set in listener configuration — calls omitting the header will be rejected.

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.165.16 NOTES ↗ RANK

Infisical is the open-source platform for secrets, certificates, and privileged access management.

> Infisical's v0.165.16 release extends the Gateway with GCP enrollment, adds ClickHouse to PAM, and tightens certificate lifecycle safety while improving secret scanning throughput and secrets UI usability.

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

01 **GCP enrollment for Infisical Gateway** NEW 40

Adds GCP enrollment support for the Infisical Gateway, enabling GCP-based infrastructure to join the gateway network.

— Names capability and platform but no setup steps v0.165.16

02 **ClickHouse support in PAM** NEW 40

Adds ClickHouse as a supported target for Privileged Access Management (PAM), enabling dynamic credential management for ClickHouse databases.

— States what's supported, no configuration detail given v0.165.16

03 **Secrets UI subfolder sync and creation flow** IMPROVED 40

Surfaces the Include All Subfolders option for secret sync in the UI, making subfolder sync behavior configurable without API access, and simplifies the resource creation menu in the secrets UI, reducing the steps needed to add new resources.

— Names UI options but not exact navigation paths v0.165.16

04 **Certificate deletion blocked until expiry** BREAKING 30

Blocks certificate deletion until the certificate has expired, preventing premature removal of active certificates.

— Clear behavior change but no mechanism or override path v0.165.16

05 **Batched secret scanning** IMPROVED 15

Adds batching to secret scanning, improving throughput for large-scale secret scan operations.

— No numbers or config to act on v0.165.16

v0.165.16 NOTES ↗ LATEST STABLE

Infisical v0.165.16 adds GCP gateway enrollment, ClickHouse PAM support, secret scanning batching, and certificate deletion blocking.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.165.16
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.165.16
```

Adds GCP enrollment support for the Infisical Gateway, enabling GCP-based infrastructure to join the gateway network.

Adds ClickHouse as a supported target for Privileged Access Management (PAM), enabling dynamic credential management for ClickHouse databases.

Blocks certificate deletion until the certificate has expired, preventing premature removal of active certificates.

Adds batching to secret scanning, improving throughput for large-scale secret scan operations.

Surfaces the `Include All Subfolders` option for secret sync in the UI, making subfolder sync behavior configurable without API access.

+1 more

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

NETRESEC

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Network forensics, packet sniffers and IT security products. Download NetworkMiner and other free software for network security analysis.

> NETRESEC introduced FlowCarp, a new statistical protocol/malware-family classifier with a trainable CLI and a free hosted demo service for quick PCAP analysis.

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

01 **FlowCarp CLI for statistical protocol and C2 classification**

NEW 91

The `flowcarp` CLI adds `--input`, `--output`, and `--model` flags to build and apply statistical protocol models from PCAP files. Custom models for proprietary or novel C2 protocols can be trained with `flowcarp --output model:<label>:<tag>:<file>.fcp` (e.g. `flowcarp --input xyz.pcap --output model:MalwareXYZ:MALICIOUS:malwarexyz.fcp`) and then applied to new captures with `flowcarp --model malwarexyz.fcp --input new.pcap --output -`. Classification distinguishes C2 traffic by packet sizes, timing, message-length distributions, request/response behavior, and protocol structure, covering 600+ protocols on the commercial license, and assigns a single protocol or malware-family label per flow rather than the multi-alert overlap common in traditional IDS-based C2 classification.

Train a custom model from sandbox-captured C2 sessions so FlowCarp can recognize a new malware family in future captures.

```
$ flowcarp --input xyz.pcap --output
model:MalwareXYZ:MALICIOUS:malwarexyz.fcp
```

Classify flows in a newly captured PCAP against a previously trained custom C2 protocol model.

```
$ flowcarp --model malwarexyz.fcp --input new.pcap --output -
```

Build a custom statistical model for a newly discovered C2 protocol from sandbox PCAPs so FlowCarp can recognize it in future captures.

```
$ flowcarp --input xyz.pcap --output
model:MalwareXYZ:MALICIOUS:malwarexyz.fcp
```

Validate a newly built C2 protocol model against captures from additional sandbox executions to check generalization before deploying.

```
$ flowcarp --model malwarexyz.fcp --input new.pcap --output -
```

— Names exact flags, model syntax, and mechanism; missing full protocol list.

changelog-20260924-5c21cf84

02 **Free FlowCarp demo web service** NEW 78

A free demo web service at `https://demo.flowcarp.com` accepts raw PCAP uploads via `curl --data-binary` and returns a JSON response listing detected flows, protocols, and alerts, letting users classify traffic without running FlowCarp locally.

Quickly identify protocols and C2 alerts in a PCAP using the free demo service without standing up a local FlowCarp instance.

```
$ curl --data-binary @dump.pcap https://demo.flowcarp.com | jq
```

— Exact URL and curl command given; no detail on rate limits or retention.

changelog-20260924-5c21cf84

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

NETRESEC: Adds `--input`, `--output`, and `--model` flags to the `flowcarp` CLI for building statistical protocol models from PCAP files and classifying traffic against them.

└──▷ TRY IT

Train a custom model from sandbox-captured C2 sessions so FlowCarp can recognize a new malware family in future captures.

```
$ flowcarp --input xyz.pcap --output
model:MalwareXYZ:MALICIOUS:malwarexyz.fcp
```

Classify flows in a newly captured PCAP against a previously trained custom C2 protocol model.

```
$ flowcarp --model malwarexyz.fcp --input new.pcap --output -
```

Quickly identify protocols and C2 alerts in a PCAP using the free demo service without standing up a local FlowCarp instance.

```
$ curl --data-binary @dump.pcap https://demo.flowcarp.com | jq
```

Adds `flowcarp --output model:<label>:<tag>:<file>.fcp` syntax to train and save a custom protocol model from example capture files, enabling detection of proprietary or novel C2 protocols without hand-written signatures.

Exposes a free demo web service at `https://demo.flowcarp.com` accepting raw PCAP uploads via `curl --data-binary` and returning a JSON response with detected flows, protocols, and alerts.

Introduces FlowCarp statistical protocol classification as a malware-family identification layer, distinguishing C2 traffic by packet sizes, timing, message-length distributions, request/response behavior, and protocol structure — covering 600+ protocols on the commercial license.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Chainguard

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

RANK

Chainguard provides trusted open source artifacts for every layer of your modern software stack—containers, language libraries, and VM images.

> Chainguard published a tutorial for using Sonatype Nexus raw repositories as pull-through caches for its apk packages, and documented a breaking default-user change in the sonar-scanner-cli image along with workarounds.

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

01 Pull-through caching via Sonatype Nexus for apk packages NEW

45

Chainguard published a tutorial explaining how to pull packages from Chainguard package repositories through Sonatype Nexus raw repositories, using them as pull-through caches for `apk` packages.

— Names the integration but no concrete steps given changelog-20260924-4ed83dfe

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Chainguard: Publishes tutorial for pulling packages from Chainguard package repositories through Sonatype Nexus raw repositories as pull-through caches for `apk` packages.

└──▷ **USE IT**

Restore the previous `sonar-scanner-cli` working-directory behavior in a Kubernetes pod spec after the October 12 default-user change.

yaml

```yaml
securityContext:
  runAsUser: 0
```

Override the scanner working directory via a JVM property so `report-task.txt` lands in the expected project path instead of `/tmp/.scannerwork`.

```
$ docker run --rm -v $(pwd):/usr/src cgr.dev/chainguard/sonar-scanner-cli \
  -Dsonar.working.directory=/usr/src/.scannerwork
```

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Tencent Cloud CubeSandbox

Feedback ⓘ v0.7.2 NOTES ↗

RANK

Instant, Concurrent, Secure & Lightweight Sandbox for AI Agents.

> CubeSandbox v0.7.2 makes cube-node run on the host network by default (a breaking change requiring an explicit confirmation flag), adds a JuiceFS storage plugin and idle-memory reclaim to pack more sandboxes per node, and introduces a local S3 cache for cross-node snapshot restores.

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

01 **Host-network mode default for cube-node** BREAKING 88

cube-node now runs on the host network by default in Kubernetes deployments, keeping sandbox network namespaces alive across Pod restarts (#1815). Switching an existing cluster to this mode requires draining compute nodes and passing the explicit `hostNetworkChangeAck` confirmation flag; without it the migration cannot proceed.

— Names exact flag and migration steps, thorough breaking-change detail v0.7.2

02 **`cubeNode.network.mtu` config key** NEW 81

Adds the `cubeNode.network.mtu` config key to pin guest MTU when the node's uplink NIC MTU cannot be auto-detected, preventing silently broken sandbox networking on deployments with constrained overlay CNIs (#1674).

Pin guest MTU to match a constrained overlay CNI uplink so sandbox networking is not silently broken on deployment.

```yaml
cubeNode:
  network:
    mtu: 1450
```

— Exact config key, example, and rationale given v0.7.2

03 **`timeout` parameter for `sandbox.connect`** NEW 70

The Python SDK's `sandbox.connect` now accepts a `timeout` parameter to control connection deadlines, letting callers cap how long a connect attempt may block (#1352).

Set a hard deadline on a sandbox connection so a hung connect attempt does not block automation pipelines indefinitely.

```python
sandbox.connect(timeout=10)
```

— Names exact parameter and example but no deeper mechanism v0.7.2

04 **Configurable template and component storage backends** IMPROVED 65

Cubelet can now configure the template install path so large templates can be moved to the data disk (#1675). Component repository and template artifact storage is extended to support S3, local disk, PVC, and NFS backends (#1749).

— Names backends and issue numbers, no exact config key given v0.7.2

THINNER COVERAGE BELOW

05 **Environment variables in `run_code` interpreter** IMPROVED 55

The Python SDK's `run_code` (E2B-compatible code-interpreter interface) can now read environment variables injected at sandbox creation (#1323).

— Names interface and behavior but no exact env var syntax v0.7.2

06 **JuiceFS storage plugin for sandbox volumes** NEW 45

Adds a JuiceFS storage plugin so sandboxes can mount a JuiceFS filesystem as a persistent volume with full POSIX semantics.

— Describes capability but no config surface given v0.7.2

07 **Local cache for cross-node S3 restores** NEW 45

Adds a whole-object local S3 cache for cross-node restore reads, so repeated reads and re-imports of the same snapshot hit the local cache instead of S3.

— Explains mechanism but no config key to invoke it v0.7.2

08 **Automatic idle guest memory reclaim** NEW 35

Adds automatic reclaim of idle guest memory back to the host, allowing a single node to host more sandboxes.

— States benefit but no mechanism or tunable exposed v0.7.2

└──▷ **BREAKING ON UPGRADE**

! cube-node now runs on the host network by default in K8s deployments; switching an existing cluster requires draining compute nodes and passing explicit `hostNetworkChangeAck` confirmation — without this, the migration cannot proceed.

v0.7.2 NOTES ↗ LATEST STABLE

CubeSandbox v0.7.2 adds JuiceFS volume plugin, idle guest memory reclaim, cross-node S3 restore cache, and `sandbox.connect` timeout parameter.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.7.2
https://github.com/TencentCloud/CubeSandbox.git
# already have the repo? check out this version:
$ git checkout v0.7.2
```

└──▷ **USE IT**

Pin guest MTU to match a constrained overlay CNI uplink so sandbox networking is not silently broken on deployment.

```yaml
cubeNode:
  network:
    mtu: 1450
```

Set a hard deadline on a sandbox connection so a hung connect attempt does not block automation pipelines indefinitely.

```python
sandbox.connect(timeout=10)
```

Adds `timeout` parameter to `sandbox.connect` to control connection deadlines (#1352).

Adds `cubeNode.network.mtu` config key to pin guest MTU when node uplink NIC MTU cannot be auto-detected (#1674).

Adds `hostNetworkChangeAck` confirmation flag required when switching an existing cluster to host-network mode for cube-node Pods (#1815).

Adds JuiceFS storage plugin so sandboxes can mount a JuiceFS filesystem as a persistent volume with full POSIX semantics.

Adds automatic reclaim of idle guest memory back to the host, allowing a single node to host more sandboxes.

+5 more

└──▷ **BREAKING ON UPGRADE**

! cube-node now runs on the host network by default in K8s deployments; switching an existing cluster requires draining compute nodes and passing explicit `hostNetworkChangeAck` confirmation — without this, the migration cannot proceed.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

The Genie One MCP server is now generally available (2026-09-25) BLOG ↗ RANK

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

> Databricks shipped GA of the Genie One MCP server (with its Beta endpoint sunsetting October 31, 2026), major Unity Catalog UDF advances letting Scala and Python functions securely access secrets and credentials, new in-SQL prediction functions, a Beta no-code app builder, and a suite of Unity Gateway tooling for centrally configuring and tracing coding agents.

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

01 **Scala UDFs can access Unity Catalog secrets** NEW 95

Unity Catalog Scala UDFs declare required secrets in a `SECRETS` clause and run under the function owner's permissions, so callers need no direct secret access. Requires `environment_version` set to `6` or above (including session-scoped UDFs on serverless compute), or on classic compute Databricks Runtime 19 or above with standard access mode; secret-accessing Scala UDFs are limited to serverless SQL warehouses among SQL warehouse types, with dedicated access mode compute unable to invoke them directly.

Define a Unity Catalog Scala UDF that reads a secret without exposing it to callers — the SECRETS clause grants owner-permission access so callers need no direct secret privilege.

```
$ CREATE OR REPLACE FUNCTION my_catalog.my_schema.fetch_api_key()
RETURNS STRING
LANGUAGE SCALA
SECRETS ('my_scope' = 'my_key')
ENVIRONMENT_VERSION 6
AS $$
  dbutils.secrets.get(scope = "my_scope", key = "my_key")
$$;
```

— Includes a full runnable CREATE FUNCTION example with exact syntax

Scala UDFs can access Unity Catalog sec…

02 **Genie One MCP server reaches general availability** BREAKING

90

The `system.ai.genie_one_mcp` MCP Service in Unity Gateway is now generally available, exposing Genie as a conversational tool over the Model Context Protocol so any MCP client or agent (Claude, ChatGPT, Cursor) can query Databricks data with Unity Catalog permissions enforced on every request. The Beta endpoint `https://<workspace-hostname>/api/2.0/mcp/genie` is deprecated and will be sunset on October 31, 2026; workloads must migrate to `system.ai.genie_one_mcp` before that date.

— Names exact endpoint and sunset date but shows no migration example

The Genie One MCP server is now general…

---

03 **Unity Catalog Python UDFs reach general availability** NEW

85

Scalar and Batch Unity Catalog Python UDFs are now GA, adding named handlers for row-at-a-time scalar UDFs, removal of the five-UDF-call-per-query limit, support for installing custom dependencies, and access to Unity Catalog service credentials and secrets from within UDF functions. Setting `environment_version` to `6` or above in the UDF definition unlocks these scalar capabilities on serverless compute, and scalar TIMESTAMP inputs are now aligned with PySpark behavior.

— Names environment_version key and limits but no code sample

Scalar and Batch Unity Catalog Python U…

---

04 **ug CLI connects coding agents to Unity Gateway** NEW

65

Adds the `ug` CLI to connect coding agents — including Claude Code, Codex, and Pi — to Unity Gateway, handling authentication, configuration, MCP server registration, and spend visibility in the terminal.

— Names CLI and agents but gives no exact command syntax

Connect coding agents to Unity Gateway …

---

05 **ai_predict_class and ai_predict_value SQL functions** NEW

60

Adds `ai_predict_class` SQL function to train a classification model and predict class labels for rows where the target column is `NULL`, and `ai_predict_value` SQL function to train a regression model and predict numeric values for rows where the target column is `NULL`.

— Names both functions and behavior but no example query shown product docs

---

06 **Coding agent activity in unified tracing** IMPROVED

60

Unity Gateway unified tracing now captures coding-agent activity — tools and skills used — alongside model requests, viewable in `Govern > Traces` or via Genie.

Review coding-agent tool and skill usage alongside model requests in the Govern traces UI.

> 📍 In the Databricks console, go to Govern › Traces to view unified traces that now include coding-agent activity.

— Gives exact UI navigation path but little other detail

Unified tracing now includes coding age…

THINNER COVERAGE BELOW

07 **No-code governed AI app building (Beta)** NEW 50

Adds Genie App Builder to author governed data and AI apps in natural language without writing code, App Spaces to enforce governance policies across apps built with Genie App Builder, and Serverless Micro Apps that scale to zero when idle — all in Beta.

— Names three Beta components but no config or usage detail

Build governed data and AI apps without…

08 **Central configuration for coding agents** NEW 50

Enables central configuration of coding agents — including models, MCP servers, skills, a default harness, and Smart Routing — with settings applied at agent launch or rolled out via mobile device management (MDM).

— Describes scope but no config keys or commands shown

Centrally configure coding agents (2026…

└──▷ BREAKING ON UPGRADE

! The Beta endpoint `https://<workspace-hostname>/api/2.0/mcp/genie` is deprecated and will be sunset on October 31, 2026; all workloads must be moved to the `system.ai.genie_one_mcp` MCP Service before that date.

The Genie One MCP server is now generally available (2026-09-25) BLOG ↗

LATEST STABLE

Genie One MCP server hits GA as `system.ai.genie_one_mcp` in Unity Gateway, exposing Genie to any MCP client with Unity Catalog permissions enforced.

└──▷ BREAKING ON UPGRADE

! The Beta endpoint `https://<workspace-hostname>/api/2.0/mcp/genie` is deprecated and will be sunset on October 31, 2026; all workloads must be moved to the `system.ai.genie_one_mcp` MCP Service before that date.

Build governed data and AI apps without code (Beta) (2026-09-23) BLOG ↗

STABLE

Databricks adds no-code governed AI app building via Genie App Builder, App Spaces, and Serverless Micro Apps (Beta).

Adds Genie App Builder to author governed data and AI apps in natural language without writing code (Beta).

Adds App Spaces to enforce governance policies across apps built with Genie App Builder (Beta).

Adds Serverless Micro Apps that scale to zero when idle (Beta).

Scala UDFs can access Unity Catalog secrets (2026-09-23) BLOG ↗ STABLE

Scala UDFs can now read Unity Catalog secrets via a SECRETS clause, with caller or owner permissions depending on UDF scope.

└──▷ **TRY IT**

Define a Unity Catalog Scala UDF that reads a secret without exposing it to callers — the SECRETS clause grants owner-permission access so callers need no direct secret privilege.

```
$ CREATE OR REPLACE FUNCTION my_catalog.my_schema.fetch_api_key()
RETURNS STRING
LANGUAGE SCALA
SECRETS ('my_scope' = 'my_key')
ENVIRONMENT_VERSION 6
AS $$
  dbutils.secrets.get(scope = "my_scope", key = "my_key")
$$;
```

Unity Catalog Scala UDFs declare required secrets in a `SECRETS` clause and run under the function owner's permissions, so callers need no direct secret access.

Requires `environment_version` set to `6` or above for Unity Catalog Scala UDFs (and session-scoped UDFs on serverless compute) to access secrets; on classic compute, requires Databricks Runtime 19 or above with standard access mode.

Limits secret-accessing Unity Catalog Scala UDFs to serverless SQL warehouses only among SQL warehouse types; dedicated access mode compute does not support direct invocation.

Scalar and Batch Unity Catalog Python UDFs are now generally available (2026-09-23)

BLOG ↗ STABLE

Unity Catalog Python UDFs reach GA with named handlers, dependency installs, and credential access — five-UDF-per-query limit removed.

Supports named handlers for row-at-a-time scalar UDFs in Unity Catalog Python UDFs.

Removes the five-UDF-call-per-query limit for Unity Catalog Python UDFs.

Enables installing custom dependencies in Unity Catalog Python UDFs.

Enables access to Unity Catalog service credentials and secrets from within Python UDF functions.

Sets `environment_version` to `6` or above in the UDF definition to unlock scalar UDF capabilities on serverless compute and other newly GA features.

`+1 more`

`Centrally configure coding agents (2026-09-22)` **BLOG ↗** STABLE

Databricks adds central configuration for coding agents, covering models, MCP servers, skills, routing, and MDM rollout.

`Connect coding agents to Unity Gateway with the ug CLI (2026-09-22)` **BLOG ↗**

STABLE

New `ug` CLI connects coding agents (Claude Code, Codex, Pi) to Unity Gateway with one command.

`Unified tracing now includes coding agent activity (2026-09-22)` **BLOG ↗**

STABLE

Unity Gateway unified tracing now captures coding-agent tool and skill activity alongside model requests.

└──▷ **HOW TO FIND IT**

Review coding-agent tool and skill usage alongside model requests in the Govern traces UI.

> 📍 In the Databricks console, go to Govern › Traces to view unified traces that now include coding-agent activity.

Adds coding-agent activity (tools and skills used) to Unity Gateway unified tracing, viewable in `Govern > Traces` or via Genie.

WAS THIS USEFUL? 👍 👎

Google SecOps

Feedback ⓘ September 21, 2026 (2026-09-21)

BLOG ↗ RANK

Meet your business challenges head on with AI and cloud computing services from Google, including security, data management, and hybrid & multi-cloud.

> Google SecOps now auto-enriches direct-ingested Google Cloud logs with an organization ID for multi-org filtering and RBAC, lets admins self-serve the Bindplane Enterprise license key, and announces two upcoming breaking changes to OAuth scopes and IOC feeds.

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

01 GCP organization ID enrichment in UDM events NEW 95

Google SecOps adds a `metadata.base_labels.ingestion_kv_labels` field with key `gcp_organization_id` to UDM events ingested via Google Cloud direct ingestion, enabling multi-organization log source identification, detection rule scoping, and Data RBAC without additional sink or feed configuration.

Scope a YARA-L detection rule to a specific Google Cloud organization in a multi-org environment by filtering on the auto-enriched org ID label.

```text
// In your YARA-L 2.0 rule, filter UDM events by org ID:
events:

$e.metadata.base_labels.ingestion_kv_labels["gcp_organization_id"
] = "123456789012"
```

— Names exact field/key and includes a runnable YARA-L filter example

September 21, 2026 (2026-09-21)

02 Legacy IOC feeds deprecated in favor of GTI_IOC DEPRECATED 75

After March 18, 2027, the `MANDIANT_ACTIVE_BREACH_IOC`, `MANDIANT_FUSION_IOC`, and `OPEN_SOURCE_INTEL_IOC` feeds will be removed; users must migrate to the `GTI_IOC` feed before that date.

— Names all deprecated feeds and replacement with exact deadline product docs

03 Self-service Bindplane Enterprise license key download NEW 70

Admins can download the Bindplane Enterprise (Google Edition) license key directly from the console via Settings > SIEM Settings > Collection Agents to onboard a new collection agent without raising a support ticket.

Download the Bindplane Enterprise license key from the console to onboard a new collection agent without raising a support ticket.

> 📍 In the Google SecOps console, go to Settings › SIEM Settings › Collection Agents, then select 'Download Bindplane Enterprise (Google Edition) license key'.

— Exact UI navigation path given, minimal further mechanism `product docs`

---

04 `chronicle.readonly` OAuth scope loses write access

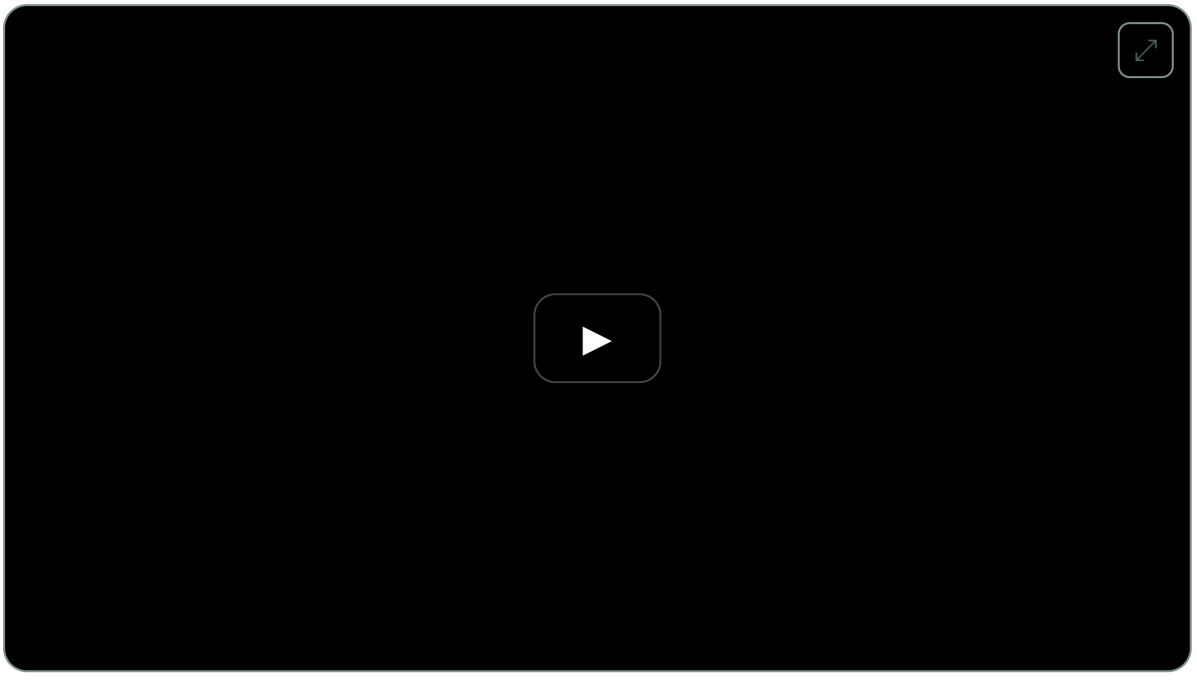
`BREAKING`

70

Effective January 25, 2027, write permissions will be removed from the `chronicle.readonly` OAuth scope; any workflow performing write operations must be updated to use the `chronicle` OAuth scope.

— Clear future date and migration target named, no further detail `product docs`

└──▷ ALSO FROM THESE RELEASES



└──▷ BREAKING ON UPGRADE

! Effective January 25, 2027, write permissions will be removed from the `chronicle.readonly` OAuth scope; any workflow performing write operations must be updated to use the `chronicle` OAuth scope.

! After March 18, 2027, the `MANDIANT_ACTIVE_BREACH_IOC`, `MANDIANT_FUSION_IOC`, and `OPEN_SOURCE_INTEL_IOC` feeds will be removed; migrate to the `GTI_IOC` feed before that date.

September 21, 2026 (2026-09-21) BLOG ↗ STABLE

Google SecOps enriches directly ingested logs with `gcp_organization_id` to simplify multi-org filtering and RBAC scoping.

Adds `gcp_organization_id` under `metadata.base_labels.ingestion_kv_labels` in UDM events for logs ingested via Google Cloud direct ingestion, enabling multi-organization filtering, detection rule scoping, and Data RBAC without additional sink or feed configuration.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Sophos Mobile

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

RANK

Sophos Mobile is an enterprise mobile device management and security solution that protects iOS and Android devices.

> Sophos Mobile expands Apple Business profile controls to let admins skip additional setup assistant screens during iOS/iPadOS enrollment.

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

01 **New setup assistant skip options in Apple Business profiles**

NEW 75

Adds Accessibility Appearance, Safety & Handling, and Liquid Glass settings to Apple Business profiles, letting admins skip those setup assistant steps during device enrollment. Note that Accessibility Appearance and Safety & Handling apply to iPhones and iPads only.

Skip the Liquid Glass, Accessibility Appearance, and Safety & Handling setup assistant screens when enrolling iPhones via Apple Business Manager.

> 📍 In Sophos Central, go to Mobile › Profiles › Apple Business profiles, open or create a profile, and enable the Accessibility Appearance, Safety & Handling, and Liquid Glass settings under the setup assistant options, then save and assign the profile to your device group.

— Names exact settings and gives UI navigation steps to enable them.

changelog-20260924-e6f43b83

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Sophos Mobile adds Accessibility Appearance, Safety & Handling, and Liquid Glass settings to Apple Business profiles.

└──▷ HOW TO FIND IT

Skip the Liquid Glass, Accessibility Appearance, and Safety & Handling setup assistant screens when enrolling iPhones via Apple Business Manager.

> 📍 In Sophos Central, go to Mobile › Profiles › Apple Business profiles, open or create a profile, and enable the Accessibility Appearance, Safety & Handling, and Liquid Glass settings under the setup assistant options, then save and assign the profile to your device group.

WAS THIS USEFUL?  

Trend Vision One

Feedback ⓘ Helm Chart 3.5.4 (2026-09-17)

BLOG ↗ RANK

Trend Vision One is Trend Micro's cloud-based XDR platform providing unified security visibility and threat detection across endpoints, networks, and cloud environments.

> Trend Vision One's Helm chart shipped two releases (3.5.3, 3.5.4) that add fine-grained node-scheduling controls, excluded-nodes billing, a new primary install source, and hardening for key components.

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

01 **Per-component toleration configuration** NEW 75

The Helm chart adds `tolerations.<component>` keys plus a `tolerations.defaults` fallback in Helm values, letting operators configure tainted-node tolerations (control-plane, GPU, spot/preemptible) for individual components rather than the whole deployment.

Upgrade the Helm chart to 3.5.4 from Public ECR, applying your tolerations and other overrides in one step.

```
$ helm upgrade \
  --values overrides.yaml \
  --namespace trendmicro-system \
  trendmicro \
  oci://public.ecr.aws/trendmicro/container-security/trendmicro-container-security
```

— Names exact Helm value keys and gives an upgrade command Helm Chart 3.5.4 (2026-09-17)

02 **Public ECR as primary Helm install source** IMPROVED 65

Helm installation now uses `oci://public.ecr.aws/trendmicro/container-security/trendmicro-container-security` as the primary installation source, with the GitHub container registry retained as the alternative.

Upgrade to 3.5.3 from the new primary public ECR source to pull the chart with resource limits and NetworkPolicy controls applied.

```
$ helm upgrade \
  --values overrides.yaml \
  --namespace trendmicro-system \
  trendmicro \
  oci://public.ecr.aws/trendmicro/container-security/trendmicro-container-security
```

— Exact OCI source named with a runnable upgrade command `Helm Chart 3.5.3 (2026-09-09)`

THINNER COVERAGE BELOW

03 **Resource limits and network hardening for admission and audit components**

IMPROVED 41

Adds resource limits for the admission-controller and policy-operator components, and restricts audit-log-collector webhook ingress via NetworkPolicy.

— Names components and mechanism but no config keys or values

`Helm Chart 3.5.3 (2026-09-09)`

04 **Excluded-nodes billing and protection** NEW 21

The chart now supports excluded-nodes billing and protection.

— Single unexplained line, no mechanism or config given `Helm Chart 3.5.4 (2026-09-17)`

`Helm Chart 3.5.4 (2026-09-17)` BLOG ↗ LATEST STABLE

Helm Chart 3.5.4 adds per-component toleration config keys, ArgoCD SCC normalization, and excluded-nodes billing support.

└──▷ **TRY IT**

Upgrade the Helm chart to 3.5.4 from Public ECR, applying your tolerations and other overrides in one step.

```
$ helm upgrade \
  --values overrides.yaml \
  --namespace trendmicro-system \
  trendmicro \
  oci://public.ecr.aws/trendmicro/container-security/trendmicro-container-security
```

Adds `tolerations.<component>` per-component keys and `tolerations.defaults` fallback in Helm values to configure tainted-node

tolerations (control-plane, GPU, spot/preemptible) for individual components.

Supports excluded-nodes billing and protection.

Helm Chart 3.5.3 (2026-09-09) BLOG ↗ STABLE

Helm Chart 3.5.3 adds resource limits for admission-controller and policy-operator and restricts audit-log-collector webhook ingress via NetworkPolicy.

└──▷ **TRY IT**

Upgrade to 3.5.3 from the new primary public ECR source to pull the chart with resource limits and NetworkPolicy controls applied.

```
$ helm upgrade \
  --values overrides.yaml \
  --namespace trendmicro-system \
  trendmicro \
  oci://public.ecr.aws/trendmicro/container-security/trendmicro-container-security
```

Updates Helm installation to use `oci://public.ecr.aws/trendmicro/container-security/trendmicro-container-security` as the primary installation source, with the GitHub container registry as the alternative.

WAS THIS USEFUL? 👍 👎

SentinelOne

Feedback ⓘ announced · Sep 24, 2026 PRESS ↗ RANK

SentinelOne's Singularity platform protects endpoints, cloud, identity and data with AI-driven detection, Purple AI and an agentic SOC layer.

> SentinelOne extended Wayfinder Threat Hunting from endpoints and identity to full multi-cloud coverage, adding continuous expert-led hunting across AWS, Azure, and GCP control planes.

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

01 **Wayfinder Threat Hunting for Cloud** NEW 75

Adds Wayfinder Threat Hunting for Cloud with continuous, expert-led hunting across AWS, Azure, and GCP control-plane activity, enabled via existing Singularity Marketplace plugins for each cloud provider. Coverage includes cloud control-plane abuse, IAM privilege escalation, unauthorized access, and data exfiltration — specifically hunts for IAM user enumeration, S3 bucket reconnaissance, root account logins, AKS cluster-admin credential access, suspicious IAM policy changes, AMI deregistration, telemetry destruction, and cross-tenant delegation changes. Curated indicators and behavioral rules are mapped to MITRE ATT&CK techniques and findings are enriched with Purple AI summaries for triage, integrating threat intelligence from both SentinelOne and Google Threat Intelligence across endpoints, identities, and cloud in a single hunting workflow.

— Names cloud providers, hunt types, MITRE mapping and enablement path via Marketplace plugins.

launch-20260924-27d9fad2

announced · Sep 24, 2026 PRESS ↗

Wayfinder Threat Hunting expands to AWS, Azure, and GCP cloud control planes with IAM, S3, and cross-tenant hunt coverage.

Adds Wayfinder Threat Hunting for Cloud with continuous, expert-led hunting across AWS, Azure, and GCP control-plane activity, enabled via existing Singularity Marketplace plugins for each cloud provider.

Covers cloud control-plane abuse, IAM privilege escalation, unauthorized access, and data exfiltration — including hunts for IAM user enumeration, S3 bucket reconnaissance, root account logins, AKS cluster-admin credential access, suspicious IAM policy changes, AMI deregistration, telemetry destruction, and cross-tenant delegation changes.

Maps curated indicators and behavioral rules to MITRE ATT&CK techniques and surfaces findings with enriched Purple AI summaries for triage.

Integrates threat intelligence from both SentinelOne and Google Threat Intelligence

> in a single hunting workflow across endpoints, identities, and cloud.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Tines lets teams put AI to work across their most sensitive systems, data, and processes, with the security and governance to do it safely.

> Tines expanded its AI usage API with new queryable dimensions and added markdown conversion functions and storyboard case management with formula support.

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

01 **Storyboard case management with formula support** NEW

16

The release added storyboard case management with formula support, per the release summary.

— Named feature only, no mechanism or UI detail provided. changelog-20260924-c9c47f7f

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Tines 43.4.0–43.3.0: AI usage API gains new dimensions, markdown conversion functions, and storyboard case management with formula support

Adds new dimensions to the AI usage API, expanding what usage data can be queried.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Splunk Security Content

Feedback ⓘ v6.7.0 NOTES ↗

RANK

Splunk Security Content

> Splunk Security Content v6.7.0 adds seven new macOS and Windows analytics targeting osascript abuse and Windows LOLBin/ClickFix techniques, alongside updates to credential-access, ESXi, and network detections, and flags three detections for removal in 6.10.0.

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

01 **Windows LOLBin and ClickFix detections** NEW 65

Four new Windows detections cover living-off-the-land and ClickFix techniques: `Windows Finger.exe Connecting to a Remote Host` identifies payload retrieval and remote connections through finger.exe, `Windows For Loop Usage Within Cmd.exe To Execute Commands` catches command execution via for `/f` loops (a common ClickFix technique), `Windows Node.exe Executing JS Script In Immediate Folder` flags JavaScript execution by node.exe from unusual directories, and `Windows Process Accessing IronLanguages Repository On GitHub` identifies processes reaching out to the IronLanguages GitHub repository.

— Four named detections with technique context, no query logic given v6.7.0

02 **macOS osascript abuse detections** NEW 62

Three new detections target osascript misuse on macOS: `Cisco NVM - Osascript Network Connection for a Long Duration` identifies long-lived network connections initiated via osascript, `MacOS Osascript Executing Interactive Shell` catches osascript spawning interactive shells, and `MacOS Osascript Executing JavaScript Code With ObjC` flags JavaScript for Automation code using the Objective-C bridge via osascript.

— Names three detections and what each targets, no SPL shown v6.7.0

THINNER COVERAGE BELOW

03 **ESXi detection improvements** IMPROVED 55

ESXi Audit Tampering and ESXi Bulk VM Termination detections were updated with improved SPL logic and threat metadata.

— Names both detections but not what SPL changes were made v6.7.0

04 **Network and web detection refinements** IMPROVED 55

DNS Kerberos Coercion, Internal Horizontal Port Scan NMAP Top 20, and Plain HTTP POST Exfiltrated Data detections were updated with better event context and reduced noise.

— Names three detections but no specifics on noise reduction mechanism v6.7.0

05 **Detections scheduled for removal in 6.10.0** DEPRECATED

50

MacOS - Re-opened Applications, Windows Process Injection into Notepad, and XMRIG Driver Loaded detections are scheduled for removal in ESCU version 6.10.0.

— Names exact detections and removal version, no migration guidance v6.7.0

06 **Credential-access and process-injection detection updates**

IMPROVED

43

Updates detections covering LSASS access and termination, credential dumping, Winlogon token manipulation, Rubeus ticket export, handle duplication, and process injection, described as improving analytic consistency and investigation context.

— Topics named but no specific detection titles or logic given v6.7.0

└──▷ BREAKING ON UPGRADE

! MacOS - Re-opened Applications, Windows Process Injection into Notepad, and XMRIG Driver Loaded detections are scheduled for removal in ESCU version 6.10.0.

v6.7.0 NOTES ↗ LATEST STABLE

Splunk Security Content v6.7.0 adds 7 new macOS and Windows analytics covering osascript, ClickFix chains, and LOLBin abuse.

└──▷ GET THIS VERSION

```
$ git clone --branch v6.7.0
https://github.com/splunk/security_content.git
# already have the repo? check out this version:
$ git checkout v6.7.0
```

Adds Cisco NVM - Osascript Network Connection for a Long Duration detection to identify long-lived network connections initiated via osascript on macOS.

Adds MacOS Osascript Executing Interactive Shell detection to catch osascript spawning interactive shells on macOS endpoints.

Adds MacOS Osascript Executing JavaScript Code With ObjC detection to identify JavaScript for Automation code using the Objective-C bridge via osascript.

Adds `Windows Finger.exe Connecting to a Remote Host` detection to identify payload retrieval and remote connections through finger.exe.

Adds `Windows For Loop Usage Within Cmd.exe To Execute Commands` detection to catch command execution via for `/f` loops, a common ClickFix technique.

+5 more

└─▷ **BREAKING ON UPGRADE**

! `MacOS - Re-opened Applications`, `Windows Process Injection into Notepad`, and `XMRIG Driver Loaded` detections are scheduled for removal in ESCU version 6.10.0.

WAS THIS USEFUL? 👍 👎

Tenzir

Feedback ⓘ

Tenzir Platform v1.40.0: Azure managed identity for Key Vault secret stores (2026-09-23)

BLOG ↗ RANK

Tenzir is a data pipeline engine for security teams that collects, transforms, enriches, and routes telemetry.

> Tenzir Platform v1.40.0 adds Azure managed identity authentication for Key Vault secret stores, JSON output for the workspace creation CLI, and the ability to force stop Explorer pipelines.

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

01 **Azure managed identity for Key Vault secrets** NEW 55

Workspace configuration can now authenticate to Key Vault secret stores using Azure managed identity, eliminating the need for long-lived client secrets.

— Explains benefit but no exact config key or setup steps given

Tenzir Platform v1.40.0: Azure managed …

02 **JSON output for workspace creation CLI** NEW 30

Adds a JSON output flag to the workspace creation CLI command.

— Flag name and exact syntax not specified Tenzir Platform v1.40.0: Azure managed …

03 **Force stop Explorer pipelines** NEW 25

Enables force stopping of Explorer pipelines.

— No mechanism or UI path given, just a bare capability name

Tenzir Platform v1.40.0: Azure managed …

> Tenzir Platform v1.40.0: Azure managed identity for Key Vault secret stores (2026-09-2
>
> BLOG ↗ LATEST STABLE
>
> Tenzir Platform v1.40.0 adds Azure managed identity auth for Key Vault secret stores and JSON output to the workspace creation CLI.
>
> Adds a JSON output flag to the workspace creation CLI command.
>
> Supports Azure managed identity authentication for Key Vault secret stores, eliminating the need for long-lived client secrets in workspace configuration.
>
> Enables force stopping of Explorer pipelines.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

runZero is an asset discovery and inventory platform that identifies and maps all devices on networks without requiring agents.

> runZero refined role-based access controls for saved queries and expanded its custom integration script library to over 100 products while opening script access to all administrators.

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

01 **Custom integration script library expanded and opened to all admins**

IMPROVED 45

The custom integration script library now covers more than 100 products and is available in the console, with individual scripts modifiable or usable as templates. Scripts are now available to all administrators, even those with narrow access scopes.

— States scope and access change but no specific script names or config

changelog-20260924-d51cf482

> changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE
>
> runZero: Custom integration script library now covers more than 100 products, available in the console with individual scripts modifiable or usable as templates.
>
> Custom roles granting Saved queries read can now view the Queries page and table without inventory access; roles granting write or manage can create, edit, copy, import, and delete saved queries.
>
> Scripts are now available to all administrators, even those with narrow access scopes.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

iDenfy

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

> iDenfy's Android SDK 9.2.0 gained landscape mode, low-light capture improvements, and UI theming, while a cross-platform Smart-ID refresh state landed alongside a breaking iOS deployment-target bump.

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

01 **Android SDK 9.2.0 UI and capture enhancements** NEW 50

Android SDK 9.2.0 adds landscape mode support, blur and glare detection for older devices, and color scheme customization with XML fallbacks.

— Names three additions but no config keys or flags changelog-20260924-928497ee

02 **Smart-ID refresh state** NEW 40

Adds a Smart-ID refresh state on Android SDK 9.2.0 and iOS SDK 9.2.3.

— States the feature and SDK versions but no mechanism changelog-20260924-928497ee

└──▷ **BREAKING ON UPGRADE**

! iOS SDK 9.2.2 raises the minimum deployment target to iOS 15.

> changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE
>
> iDenfy: Adds landscape mode support on Android SDK 9.2.0.
>
> Adds blur and glare detection for older devices on Android SDK 9.2.0.
>
> Adds color scheme customization with XML fallbacks on Android SDK 9.2.0.
>
> Adds Smart-ID refresh state on Android SDK 9.2.0 and iOS SDK 9.2.3.
>
> └──▷ **BREAKING ON UPGRADE**
>
> ! iOS SDK 9.2.2 raises the minimum deployment target to iOS 15.

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

RANK

Govern workforce access, deploy AI agents, and get work done with C1.

> ConductorOne shipped a single guided setup flow for creating apps and connectors, added search and faceted filtering to the automation builder, and introduced a Transformation rule step for automatic assignee routing — alongside three breaking API changes to revoke criteria matching, secret deletion, and task spend-remedy handling.

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

01 **Search and filtering in automation builder** IMPROVED 60

The automation builder now supports searchable step and trigger-type search by name or description, and faceted, searchable resource-type pickers with real counts have been added to `Grant Found`, `Grant Deleted`, and revoke criteria.

— Named areas of change but no exact navigation or command. changelog-20260924-e7a9b8d7

THINNER COVERAGE BELOW

02 **Guided setup flow for apps and connectors** NEW 30

Adds a single guided setup flow for creating apps and adding connectors.

— States the feature but no mechanism or UI path given. changelog-20260924-e7a9b8d7

└──▷ **ALSO FROM THESE RELEASES**

. . . .

└──▷ **BREAKING ON UPGRADE**

! A revoke criteria scoped only by a legacy resource-type name now fails closed (fails the step) instead of silently matching every app sharing that name — re-scope any automation still relying on name-only matching.

! `DELETE /api/v1/secrets/{vault_id}` now deletes only the secret and no longer also revokes its credential; use the new `POST /api/v1/app-secrets-admin/revoke` to queue a credential revoke separately. Service SDK `Revoke` method is renamed to `Delete`.

! `GET /api/v1/tasks/{task_id}/spend-remedy-status` has been removed — migrate to the task APIs. JS/TS clients must handle spend-control revision values as strings, and integrations with a closed list of task action types must handle the new `TYPE_SPEND_REMEDY = 7`.

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

ConductorOne: Adds a single guided setup flow for creating apps and adding connectors.

└──▷ HOW TO FIND IT

Automatically route findings to the right owner in a Transformation rule instead of relying on manual assignment.

📍 In the C1 console, go to Automation › Transformation rules, open or create a rule, add a Set assignee step, and select the target user — this assignment takes precedence over any later manual change.

Adds searchable step and trigger-type search by name or description in the automation builder.

Adds faceted, searchable resource-type pickers with real counts to Grant Found, Grant Deleted, and revoke criteria.

└──▷ BREAKING ON UPGRADE

! A revoke criteria scoped only by a legacy resource-type name now fails closed (fails the step) instead of silently matching every app sharing that name — re-scope any automation still relying on name-only matching.

! `DELETE /api/v1/secrets/{vault_id}` now deletes only the secret and no longer also revokes its credential; use the new `POST /api/v1/app-secrets-admin/revoke` to queue a credential revoke separately. Service SDK `Revoke` method is renamed to `Delete`.

! `GET /api/v1/tasks/{task_id}/spend-remedy-status` has been removed — migrate to the task APIs. JS/TS clients must handle spend-control revision values as strings, and integrations with a closed list of task action types must handle the new `TYPE_SPEND_REMEDY = 7`.

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

The easiest way to add authentication and user management to your application. Purpose-built for React, Next.js, Remix, and “The Modern Web”.

> Clerk added CIMD support for public OAuth clients like MCP and introduced a dedicated SMS logs view for tracing verification code delivery.

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

01 **SMS logs view in Application Logs** NEW 88

A new SMS logs view in Application Logs tracks five delivery event types — `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` — for phone verification codes, OTPs, and password reset codes. All events for a single SMS message are grouped under a shared trace, enabling end-to-end timeline reconstruction for a single recipient; failures can be filtered by destination phone number and inspected via a `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`).

Investigate why a user did not receive their verification code by filtering SMS delivery events to a single phone number.

> 📍 In the Clerk Dashboard, open the SMS logs view under Application Logs and filter by the destination phone number. Check for `sms.failed` or `sms.undeliverable` events and inspect the `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) to determine the corrective action.

— Names all event types and reason codes with clear UI path changelog-20260924-3af810e3

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Clerk: Enables Publish CIMD support in OAuth application settings so MCP and other public OAuth clients can identify themselves via an HTTPS URL as their `client_id`, without a pre-issued client secret or Dynamic Client Registration.

└──▷ **HOW TO FIND IT**

Lock your OAuth server to only clients you have explicitly reviewed, preventing unknown public clients from initiating flows.

> 📍 In the Clerk Dashboard, go to OAuth applications settings › Client onboarding. Enable Publish CIMD support, then set Client admission to Pre-registered clients only and add trusted clients from the Applications tab.

Investigate why a user did not receive their verification code by filtering SMS delivery events to a single phone number.

> 📍 In the Clerk Dashboard, open the SMS logs view under Application Logs and filter by the destination phone number. Check for `sms.failed` or `sms.undeliverable` events and inspect the `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) to determine the corrective action.

Adds an SMS logs view in Application Logs with five delivery event types — `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` — for phone verification codes, OTPs, and password reset codes.

Groups all events for a single SMS message under a shared trace, enabling end-to-end timeline reconstruction for a single recipient.

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Secure users, AI agents, and more with Auth0, an easy-to-implement, scalable, and adaptable authentication and authorization platform.

> Auth0's biggest changes this window are a new Action method for assigning custom user IDs at signup, real-time enterprise connection events, reusable Action Modules, and a My Organization API for B2B member management, alongside early-access B2B Connect improvements.

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

01 **Custom user IDs at signup via Actions** NEW 85

New api.user.setUserId() method in `onExecutePreUserRegistration` Actions lets you assign custom user IDs (UUIDs, domain-specific, or legacy-compatible formats) to new database-connection users at signup, replacing the need for self-hosted signup pages or bulk import.

Assign a UUID-based custom user ID at signup to keep identity consistent with your existing systems, instead of accepting Auth0's auto-generated `auth0|<id>` format.

javascript

```
exports.onExecutePreUserRegistration = async (event, api) => {
  const crypto = require('crypto');
  api.user.setUserId(crypto.randomUUID());
};
```

— Names the exact method and includes a runnable code example.

changelog-20260924-66040154

02 **B2B member management via My Organization API** NEW 85

Adds the My Organization API and Embeddable UI Components for B2B member management, supporting bulk member invitations (up to 10 per request), role assignment, and a four-tier access level model (`none` , `readonly` , `limited` , `full`). Includes permission-aware UI components that show or hide actions based on the authenticated user's effective permissions surfaced from the ID token with no additional API call, plus tenant log events for all member, invitation, and role operations for B2B observability.

Activate the My Organization API to enable self-service B2B member management for organization administrators.

📍 In the Auth0 Dashboard, go to Applications › APIs, locate the My Organization API, and activate it for your tenant.

— Names the API, access tiers, limits, and dashboard activation steps.

changelog-20260924-66040154

03 **Action Modules for shared Action code** NEW 70

Adds Action Modules, enabling reuse of custom functions across multiple Actions and Triggers, each with independent secrets and dependencies. Adds unit testing support for Action Modules via helpers in the `@auth0/actions` package, plus IntelliSense autocompletion and type hints for Action Modules in the Auth0 Dashboard and external editors via `@auth0/actions`.

— Names the package and capabilities but gives no concrete command.

changelog-20260924-66040154

04 **Real-time enterprise connection events** NEW 68

Adds `connection.created`, `connection.updated`, and `connection.deleted` events for enterprise connections (Azure AD, ADFS, Google Workspace, Active Directory, LDAP, Ping), deliverable to a webhook or Amazon EventBridge as CloudEvents-formatted JSON.

— Names event types and delivery channels but no setup steps. changelog-20260924-66040154

THINNER COVERAGE BELOW

05 **B2B Connect early access enhancements** IMPROVED 50

Extends B2B Connect (Early Access) to support automatic session termination when an enterprise IdP ends a session, email-domain-based enterprise connection discovery, and direct application connections to Auth0 without an external authorization server.

— Lists new capabilities but no configuration or usage detail.

changelog-20260924-66040154

└──▷ ALSO FROM THESE RELEASES

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Auth0 adds Action Modules, B2B Member Management API, custom user IDs at signup, and real-time enterprise connection events

└──▷ USE IT

Assign a UUID-based custom user ID at signup to keep identity consistent with your existing systems, instead of accepting Auth0's auto-generated `auth0|<id>` format.

```javascript
exports.onExecutePreUserRegistration = async (event, api) => {
  const crypto = require('crypto');
  api.user.setUserId(crypto.randomUUID());
};
```

Activate the My Organization API to enable self-service B2B member management for organization administrators.

> 📍 In the Auth0 Dashboard, go to Applications › APIs, locate the My Organization API, and activate it for your tenant.

Adds api.user.setUserId() method in `onExecutePreUserRegistration` Actions to assign custom user IDs (UUIDs, domain-specific, or legacy-compatible formats) to new database-connection users at signup, replacing the need for self-hosted signup pages or bulk import.

Adds `connection.created`, `connection.updated`, and `connection.deleted` events for enterprise connections (Azure AD, ADFS, Google Workspace, Active Directory, LDAP, Ping), deliverable to a webhook or Amazon EventBridge as CloudEvents-formatted JSON.

Adds Action Modules, enabling reuse of custom functions across multiple Actions and Triggers, with independent secrets and dependencies per module.

Adds unit testing support for Action Modules via helpers in the `@auth0/actions` package.

Adds IntelliSense autocompletion and type hints for Action Modules in the Auth0 Dashboard and external editors via `@auth0/actions`.

+4 more

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ changelog · Sep 24, 2026

CHANGELOG ↗ RANK

Discover every SaaS and AI asset, harden security posture, and scale identity governance.

> Nudge Security shipped a bulk nudge for enforcing approved password managers and a new device code phishing detection feature that warns employees in-browser across major OAuth providers.

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

01 **Bulk nudge for approved password manager use** NEW 82

A new Use approved password manager bulk nudge lets admins filter the accounts inventory by password manager — including unapproved or unrecognized ones — multi-select the matching accounts, and send a single batched nudge per employee listing all affected accounts, rather than one nudge per account.

Sweep out unapproved password manager use across your org in one pass — employees get a single batched nudge instead of per-account noise.

> 📍 In the Nudge Security console, go to Accounts Inventory › Filter by Password Manager › select any unapproved or unrecognized password manager › multi-select all matching accounts › Send Nudge › Use approved password manager .

— Full workflow with exact UI navigation steps given changelog-20260924-13ee791c

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Nudge Security: Adds a Use approved password manager bulk nudge: filter accounts inventory by password manager (including unapproved ones), multi-select, and send one batched nudge per employee listing all affected accounts.

└──▷ **HOW TO FIND IT**

Sweep out unapproved password manager use across your org in one pass — employees get a single batched nudge instead of per-account noise.

> 📍 In the Nudge Security console, go to Accounts Inventory › Filter by Password Manager › select any unapproved or unrecognized password manager › multi-select all matching accounts › Send Nudge › Use approved password manager .

Enable device code phishing protection so employees are warned the moment they land on an OAuth device code authorization page.

> 📍 In the Nudge Security console, go to the device code phishing detection settings, enable detection, and turn on the in-browser nudge to activate warnings across Microsoft, Google, GitHub, GitLab, Salesforce, AWS, Docker, Zoom, Vercel, and OpenAI device code pages.

WAS THIS USEFUL? 👍 👎

Ping Identity

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

RANK

Read docs, explore use cases, learn best practices

> Ping Identity's rapid channel adds new CAPTCHA and QR-code options for authentication nodes, renames ThreatMetrix nodes to LexisNexis DDP, and reorganizes the admin console navigation.

└──▷ **WHAT SHIPPED · 7 FEATURES** most completely described first ? what's the number?

01 **Configurable userAgent for external REST calls** NEW

72

Adds a `userAgent` property to external REST service configuration, defaulting to `PingIdentity` for outbound requests, and overridable on a per-request basis.

— Named config key and default given, no further mechanism detail

changelog-20260924-fb4a2526

02 **JavaScript-free QR codes across registration nodes** NEW

60

Adds an option to the Combined MFA Registration, OATH Registration, and Push Registration nodes to render QR codes without using JavaScript.

— Names all three affected nodes but not the option's config path

changelog-20260924-fb4a2526

THINNER COVERAGE BELOW

03 **Dynamic radio-button choices in Choice Collector node** NEW

57

The v2.0 Choice Collector node gains the ability to make choice callbacks with dynamically generated options that render as radio-button groups instead of drop-down lists.

— Describes behaviour change with node name but no config example

changelog-20260924-fb4a2526

04 **LexisNexis DDP and IIDQA node updates** IMPROVED

55

Renames ThreatMetrix nodes to LexisNexis Dynamic Decision Platform (DDP) nodes, and updates LexisNexis InstantID Question and Answer (IIDQA) nodes to store decision API responses in shared state.

— Names both node families and the rename/behaviour change changelog-20260924-fb4a2526

05 **Managed user updates with attribute-conflict override**

IMPROVED 50

Enables updating a managed user even when the change conflicts with generated attribute values, propagating updates to all associated workforce application accounts during queued synchronization.

— Explains mechanism and scope but no config surface named changelog-20260924-fb4a2526

06 **Friendly Captcha support in CAPTCHA node** NEW 45

Adds Friendly Captcha support to the v2.0 CAPTCHA node as a privacy-focused CAPTCHA provider alternative.

— Names the node and provider but no configuration detail changelog-20260924-fb4a2526

07 **Reorganized admin console navigation** IMPROVED 35

Reorganizes the admin console left menu into logical, collapsible groups with updated labels.

— Clear UI change but no specifics on new grouping or labels changelog-20260924-fb4a2526

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Ping Identity adds Friendly Captcha support, JavaScript-free QR codes, dynamic radio-button choices, and a configurable `userAgent` for external REST calls.

Adds `userAgent` property to external REST service configuration, defaulting to `PingIdentity` for outbound requests, overridable per-request.

Adds Friendly Captcha support to the v2.0 CAPTCHA node as a privacy-focused CAPTCHA provider alternative.

Adds an option to the Combined MFA Registration, OATH Registration, and Push Registration nodes to render QR codes without using JavaScript.

Adds the ability in the v2.0 Choice Collector node to make choice callbacks with dynamically generated options that render as radio-button groups instead of drop-down lists.

Renames ThreatMetrix nodes to LexisNexis Dynamic Decision Platform (DDP) nodes, and updates LexisNexis InstantID Question and Answer (IIDQA) nodes to store decision API responses in shared state.

+2 more

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ 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.

> JumpCloud's biggest window shipped Password Vault across Admin Portal, User Portal, mobile app and browser extension, a unified cross-platform Patch Management dashboard, AI Gateway/MCP Server management with a GA MCP Server, Access Requests, non-human Agent Identities, direct SSH access in PAM, a major Active Directory Integration overhaul (agent HA, credential hardening, new Directory Insights events), and broad expansion of Android/Windows MDM, RADIUS, OIDC, SaaS Management and the Password Manager desktop app, alongside several breaking API and endpoint changes.

└──▷ **WHAT SHIPPED · 60 FEATURES** most completely described first ? what's the number?

01 **Patch Management dashboard for Windows and Apple** NEW

91

A unified Patch Management dashboard in the Admin Portal adds `Overview`, `Windows`, and `Apple` tabs for cross-platform patch compliance. The Windows dashboard offers KB-centric controls to search/filter KBs, review CVE details, approve awaiting updates, retry failed installations, and export the KB list as CSV or JSON; the Apple Compliance Dashboard tracks macOS, iOS, and iPadOS rollout progress against active update policies with filtering by platform, release date, severity, and device status. A Device Patch Report (Windows) lets admins preview up to 500 rows and download the full organization result set as CSV.

Surface Windows KB patch gaps across your fleet — filter the dashboard to Critical severity, then bulk-install missing patches immediately instead of waiting for the next policy window.

> 📍 In the Admin Portal, go to Patch Management › Windows, set the Global severity filter to `Critical`, select non-compliant devices, and choose `Install missing critical patches` › `Install immediately`.

Audit Windows patch compliance across all managed devices without any manual data collection — download the full result set for offline analysis or ticketing.

> 📍 In the Admin Portal, navigate to JumpCloud Reports › Device Patch Report (Windows). Preview up to 500 rows to spot devices with `Failed` or `Pending` KB status, then click Download to export the complete organization result set as CSV.

— Names dashboard tabs, controls, and report export path `changelog-20260924-f4d0af19`

`Release Notes 2026 (2026-09-24)`

02 **AI Gateway and MCP Server management** NEW

90

The JumpCloud MCP Server reached GA, letting compatible AI tools and code editors interact with organizational data and admin capabilities via the Model Context Protocol. A new `Access > AI & SaaS Management > Applications` view adds an `MCP Connection` column showing `Discovered`, `Connected`, or `Not Available` status per app plus an `Add MCP Server` action to connect discovered servers to AI Gateway, and a new `Dashboard` tab in AI Gateway shows org-wide MCP server usage, tool call volume, error rates, summary cards, trend charts, and per-server drill-down metrics.

Surface AI agents deployed in your environment and connect a newly discovered MCP server to AI Gateway for centralized access control.

> 📍 In the Admin Portal, go to Access › AI & SaaS Management › Applications. Check the MCP Connection column for any app showing `Discovered`, click the status,

review the listed MCP servers, then click Add MCP Server to register it in AI Gateway.

— Names exact UI columns, statuses, and dashboard metrics Release Notes 2026 (2026-09-24) Release Notes 2025 (2026-09-24)

03 **Password Vault launch across Admin, User, mobile and browser**

NEW 90

Adds Password Vault across the platform: in the Admin Portal, admins enable the vault, assign access to users and user groups, create and share credentials/websites with Connect, View, Edit, and Manage permissions, and view audit logs in Directory Insights; in the User Portal (left navigation), users manage and share credentials/websites, bulk-import via CSV or Excel, and organize items in personal and shared folders; in the JumpCloud Protect mobile app users can view secrets, autofill credentials into mobile apps, and launch websites with autofill, with TOTP codes remaining stored locally; and the JumpCloud Go browser extension gains autofill support for launch-and-autofill of non-SSO passwords and 2FA codes, secret viewing, and save-on-login/update/signup without requiring device registration.

— Names permission levels, import formats and surfaces Release Notes 2026 (2026-09-24)

04 **Remote Assist capability expansion** IMPROVED 90

Launches JumpCloud Remote Assist for free, unlimited remote desktop access and support on Windows and macOS (requiring JumpCloud Agent v1.61.0+), later adding Background Tools for remote command line and file transfer without disrupting active sessions, support for viewing/interacting with the initial Windows Server login screen, user switching between accounts during sessions, a Ctrl+Alt+Del shortcut in the Session Options toolbar, architecture-specific packages (ARM64 for Apple Silicon, x86_64 for Intel Macs), and automatic detection and termination of existing Remote Assist agent instances when an admin initiates a connection.

Remotely troubleshoot an end-user machine during an incident without interrupting the user's active session, using Background Tools command line.

📍 In the Admin Portal, go to Devices › select the target device › Background Tools › Remote Command Line, then run your diagnostic commands without launching a visible Remote Assist session.

— Names every capability and the exact agent version gate Release Notes 2024 (2026-09-24) Release Notes 2023 (2026-09-24) Release Notes 2022 (2026-09-24)

05 **ADI agent deployment flexibility and high availability**

IMPROVED 87

The AD sync agent gains a new JSON configuration file with settings for domain controller address and a flag allowing insecure LDAP connections, plus high availability: multiple sync agents can be installed with one designated primary by the ADI service, and if the primary becomes unavailable another active agent is automatically promoted. The sync agent can now be installed independently of the AD import agent, and both the import and sync agents can be installed on one or more member servers instead of on all Domain Controllers, reducing required agent deployments. The AD import agent installer now detects existing installations and updates in place instead of requiring uninstall/reinstall, and JumpCloud added official ADI support for Windows Server 2022 Domain Controllers and member servers without workarounds.

— Names config file fields and every deployment change `Release Notes 2023 (2026-09-24)`

---

06 **Disk-encryption Conditional Access condition** NEW 86

Adds a disk-encryption device condition to Conditional Access Policies, allowing or denying access based on FileVault (macOS), system drive encryption/BitLocker (Windows), or root disk encryption (Linux); admins add the Device Condition, select `Disk Encryption`, set it to `disabled`, and set the policy action to `Deny` to block unencrypted devices.

Block device access via a Conditional Access Policy when disk encryption is not enabled, covering FileVault (macOS), BitLocker (Windows), and root disk encryption (Linux).

> 📍 In the Admin Portal, go to Security › Conditional Access › New Policy, add a Device Condition, select `Disk Encryption`, set the condition to `disabled`, and set the policy action to `Deny`.

— Exact policy steps and per-OS encryption mechanisms `Release Notes 2022 (2026-09-24)`

---

07 **Dynamic Groups for devices and users** NEW 86

Adds Dynamic Groups for both devices and users, automating group membership via attribute-based rules; new device attributes include Vendor and Public IP Address, and a new user attribute Manager is available. Example: create a device group with rules like Vendor = `Apple` and OS Major Version = `14` to automatically update membership.

Automate device group membership so newly enrolled corporate laptops running a specific OS version are added to the right policy group without manual review.

> 📍 In the Admin Portal, go to Device Groups › select or create a group › enable Dynamic › add attribute rules (e.g., Vendor = `Apple`, OS Major Version = `14`) › save to activate automatic membership updates.

— Names new attributes and a concrete rule example `Release Notes 2023 (2026-09-24)`

---

08 **Private Repository application version updates** NEW 83

`Update Version` action on an existing application's `Details` tab in the Private Repository lets admins push an updated binary without re-creating the application or reassigning devices; for EXE uploads it also captures an install success exit code.

Push an updated application binary to already-enrolled devices without reassigning device groups or recreating the app entry in the Private Repository.

> 📍 In the Admin Portal, go to the application's Details tab in the Private Repository, click `Update Version`, upload the new file via the `Upload File` flow, enter an install success exit code (for EXE), and confirm deployment.

— Exact UI path and mechanism given, no numeric limits `changelog-20260924-f4d0af19`

---

09 **AI Assistant knowledge base grounding** NEW 83

AI Assistant supports an org-scoped knowledge base upload of internal PDF, TXT, DOCX, or MD files (up to 10 MB each, 25 files maximum per org) via AI Assistant > Knowledge Base, with chat answers attributed to the source document.

Ground AI Assistant in org-specific runbooks and SOPs so it can answer questions about internal processes, with answers attributed to the source document.

> 📍 In the Admin Portal, navigate to AI Assistant › Knowledge Base, click `Upload`, and add your internal PDF, TXT, DOCX, or MD files (up to 10 MB each, 25 files per org).

— Exact file types, size and count limits, and UI path `changelog-20260924-f4d0af19` `Release Notes 2026 (2026-09-24)`

---

10 **API filter escape character for pipe values** IMPROVED 81

Adds support for escaping the pipe `|` character with `\` inside `in` and `nin` API filter operator search terms, e.g. `GET /systemusers?filter=department:$in:Finance | Accounting|IT` for filtering values containing a literal pipe.

Filter users whose department is `Finance | Accounting` or `IT` using the newly escapable pipe character in API filter search terms.

```
$ curl -s -X GET 'https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT' \
  -H 'x-api-key: <your-api-key>' \
  -H 'Content-Type: application/json'
```

— Runnable filter syntax with escape character shown `Release Notes 2024 (2026-09-24)`

---

11 **Systemusers reactivate endpoint migration**

BREAKING

81

Replaces the deprecated `GET api/systemusers/reactivate/{id}` endpoint (which now returns 404) with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}`, which takes an `activation_email` field in the request body.

Migrate an automation that reactivates a user account from the deprecated GET endpoint to the new POST endpoint.

```
$ curl -s -X POST
'https://console.jumpcloud.com/systemusers/reactivate/{id}' \
  -H 'x-api-key: <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d '{"activation_email": "user@example.com"}'
```

— Runnable request with method, path and body field Release Notes 2024 (2026-09-24)

12 **ADI credential and file security hardening** IMPROVED

77

Moves AD Import Service Account credentials (username and password) and the ADI API key out of the AD import agent configuration file and into the Windows registry, masks the connect key value in the registry when the JumpCloud AD Integration Sync service starts, and consolidates all AD import files/configs under `C:\Program Files\JumpCloud\AD Integration\JumpCloud AD Import` and all AD sync files/logs/config under `C:\Program Files\JumpCloud\AD Integration\JumpCloud AD Sync`.

— Names exact registry moves and consolidated file paths product docs

13 **JumpCloud Go passwordless authentication** NEW

76

Adds the `JumpCloud Go` browser extension for hardware-protected, phishing-resistant passwordless login to JumpCloud-protected web resources on managed macOS and Windows devices via Apple Touch ID and Windows Hello, set as the default MFA method for conditional access policies on SSO apps with user and device verification on every session; a User Portal > Security > Multi-Factor Authentication section lets users manage JumpCloud Go sessions on their devices, and sessions automatically refresh when a user logs in to or unlocks their managed device.

— Names auth mechanism, portal path and refresh trigger Release Notes 2024 (2026-09-24) Release Notes 2023 (2026-09-24)

14 **Android device management expansion** NEW

75

JumpCloud expanded Android MDM across the window: zero-touch enrollment for fully managed, dedicated, and company-owned devices with a work profile; EMM enrollment and management for company-owned and personal Android devices using work profiles; a Phase 2 rollout of fully managed/dedicated device support adding 7 policies including Factory Reset Protection, Kiosk Mode, and System Updates; new EMM policies for Display Settings, Autofill (Android 8.0+), Work Account (Google enterprise account enforcement), and Default App and Intent Handling; and full Android 15 support.

— Names every policy and enrollment mode added `changelog-20260924-f4d0af19`

`Release Notes 2024 (2026-09-24)` `Release Notes 2023 (2026-09-24)`

---

15 **Network authentication protocol hardening** IMPROVED

75

Adds tracking of the source IP of incoming LDAP requests for security and traceability; adds faster EAP reconnect using cached session credentials to skip the full security handshake for recently authenticated devices; and enforces only modern TLS cipher suites for LDAP TLS connections (`TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256` , `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256` , `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384` , `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384` , `TLS_RSA_WITH_AES_128_GCM_SHA256`), removing CBC-mode and SHA-1/SHA-256 suites.

— Names every cipher suite and mechanism, no migration steps

`Release Notes 2025 (2026-09-24)`

---

16 **Forced re-authentication and AMR claims for SSO** NEW

73

Adds a `sso_auth` Directory Insights event capturing forced re-authentication events; adds `ForceAuthn='true'` (SAML) and `prompt=login` (OIDC) support so service providers can trigger step-up authentication without ending the user's existing session; and includes Authentication Methods References (AMR) claims in both SAML assertions and OIDC ID tokens to surface proof of MFA factors used at the point of access.

— Names event and exact SAML/OIDC parameters `Release Notes 2025 (2026-09-24)`

---

17 **Dynamic Groups API operator deprecation** BREAKING

72

Adds `$lte` and `$gte` filter operators for Dynamic Groups API `memberQuery` filters, replacing the deprecated `$le` and `$ge` operators; existing API scripts using `$le` / `$ge` must be updated.

— Exact deprecated and replacement operator names given `Release Notes 2024 (2026-09-24)`

---

18 **New ADI Directory Insights events** IMPROVED

71

Adds new Directory Insights events for the ADI integration: `user_update_password_provision` (logged by the AD sync agent on password attribute changes/attempts in AD), `activedirectory_agent_activate` (logged when an AD sync agent comes online), and `idsource_create` (logged when an import agent comes online); also updates the existing `activedirectory_agent_create` event with more information in the event listing and summary.

— Names every event by exact key `product docs`

---

19 **Asset Management enhancements** NEW 69

Asset Management gains file upload support (via the Invoice field under Purchase Details, plus custom fields of the File Upload type), a Disk Size field in Hardware Specifications, a Local Device Accounts display on the Asset Detail page, a dedicated Asset Manager admin role, and asset cloning that automatically clears unique fields to prevent duplicates.

— Names every field, role and mechanism added `Release Notes 2025 (2026-09-24)`

---

20 **Patch and policy compliance reporting** NEW 69

Adds a universal Browser Patch Management policy to keep Google Chrome up to date across macOS, Windows, and Linux from a single policy; a Linux Forbidden Services Policy to remove unsupported service types from Linux fleets; an OS Patch Management Policy Report showing which policies are applied to each device and current compliance status; and a Browser Patch Management Policy Report for compliance audit evidence covering devices' current browser versions.

— Names each policy and report by name `Release Notes 2022 (2026-09-24)`

---

21 **Commands: templates, queue visibility, and ad hoc execution** NEW 68

Adds a Command After Agent Install command type so admins can run enrollment commands immediately after agent install on new devices; a Running column in the commands queue to distinguish queued from actively running commands; the ability to run ad hoc commands on devices directly from the Multi-Tenant Portal and Admin Portal; and Custom Admin Command Templates for MSP admins to save and reuse commands across tenant organizations in the Multi-Tenant Portal.

— Names each command surface, no single deep mechanism `Release Notes 2023 (2026-09-24)` `Release Notes 2022 (2026-09-24)`

---

22 **Windows MDM CSP policy expansion** NEW 68

Adds six new Windows MDM policies via the Configuration Service Provider (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 separate policy to configure Simple Certificate Enrollment Protocol (`SCEP`) for Windows devices.

— Names all seven policies added, no config walkthrough Release Notes 2024 (2026-09-24)

23 **Device Monitoring and Alerting** NEW 68

Adds Device Monitoring and Alerting for Windows, macOS, and Linux endpoints, monitoring availability, software, uptime, disk, battery, policy/command/software management results, user addition, and privilege elevation events, with custom alerts via command execution monitoring.

— Names every monitored signal, no threshold/config shown Release Notes 2024 (2026-09-24)

24 **Support case portal: Case Portal to Customer Portal** IMPROVED 67

JumpCloud Case Portal launched to centralize support cases and ideas from all administrators in a filterable, clickable list; it was later replaced by a new Customer Portal with sortable columns for Last Modified Date, Account, and Requester, improved search, filter, attachment uploads, editable ticket fields (System ID, steps to reproduce), and extended access so MTP admins can view and create support cases for all managed tenant organizations.

— Names sortable fields and editable ticket fields Release Notes 2025 (2026-09-24) Release Notes 2022 (2026-09-24)

25 **Password Manager admin install/update via command template** IMPROVED 67

Admins can install and update the Password Manager desktop app via a new command template v.2.0.0 on Windows and macOS, eliminating the need for users to have admin permissions for subsequent updates.

— Names the template version and permission change

What's New: Password Manager Desktop Ap…

26 **SaaS Management launch** NEW 65

Launches SaaS Management to discover, track, categorize, warn, and block SaaS app access on user devices using the JumpCloud Go browser extension, integrated with existing JumpCloud SSO apps; admins can also manually add applications for tracking, warning, and blocking before automatic discovery, and discovery of unmanaged apps works via the browser extension alone without requiring JumpCloud Go.

— Explains mechanism and manual/automatic discovery paths Release Notes 2024 (2026-09-24)

27 **Homepage widgets for certificates, tenant reports, and trials** IMPROVED 65

Adds an MDM Upcoming Certificate Expirations widget to the Admin Portal and Multi-Tenant Portal homepages, later expanded into an MDM Certificate & Token Status Board

showing expiration status for MDM, SSO, and AzureAD tokens/certificates with an aggregate view for MSP admins; adds a Multi-Tenant Portal homepage widget showing aggregate reports available across all tenant organizations in the last 24 hours with bulk export to JSON or CSV; and adds a trial recap summary giving insight into feature usage during JumpCloud trials.

— Names widgets and export formats, no config steps Release Notes 2023 (2026-09-24) Release Notes 2022 (2026-09-24)

28 **Staged users no longer auto-provisioned to AD** BREAKING

65

Staged users in JumpCloud are no longer provisioned to Active Directory — users that were previously created in AD from the staged state will no longer be created there until their state is changed to active.

— Clear before/after behaviour change stated product docs

29 **App Catalog for Windows** NEW

65

Adds the JumpCloud App Catalog for Windows with ten pre-built apps (1Password, Adobe Acrobat Reader DC, Chrome, Firefox, Google Drive, OneDrive, OpenVPN, Slack, Webex, Zoom) including auto-update support.

— Names all ten catalog apps, no deployment steps Release Notes 2024 (2026-09-24)

30 **`.appimage` Linux build deprecated for Password Manager**

62

BREAKING

The `.appimage` Linux build of Password Manager is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

— Exact deprecation date and required migration path

What's New: Password Manager Desktop Ap…

31 **RADIUS authentication expansion** NEW

62

Adds certificate (passwordless) based authentication to JumpCloud RADIUS, supporting certificates from any X.509 compliant certificate authority, and adds RADIUS authentication with Azure AD as the identity provider, enabling RADIUS access without managing users and passwords outside of Azure AD.

— Names cert standard and IdP but no setup steps Release Notes 2022 (2026-09-24)

32 **macOS device management policies** NEW

62

Adds a Managed Login Items policy for macOS 13 Ventura to prevent users from disabling persistent software such as the JumpCloud Agent, and macOS software update

scheduling for individual devices from the Admin Portal, including forced updates and user delay/cancel controls.

— Names the macOS version and update control options Release Notes 2022 (2026-09-24)

33 **Portal navigation and login experience** IMPROVED 61

Adds an Identity First Login flow for the User Portal and SSO applications that splits username and password entry across two screens; splits the User Portal Applications area into All Applications and Favorites tabs with starring and responsive search; and adds a `⌘/Ctrl + K` keyboard shortcut for fast navigation across the Admin Portal.

— Names each UI surface, limited mechanism depth Release Notes 2024 (2026-09-24) Release Notes 2023 (2026-09-24) Release Notes 2022 (2026-09-24)

34 **Directory Insights filter refinement** IMPROVED 60

Refines the Directory Insights Application Name filter into sub-categories (Application Name , Application Display Name , Application Display Label) and Resource Name into (Resource Name , Resource Display Name , Resource Display Label).

— Names exact filter sub-categories, no usage example Release Notes 2025 (2026-09-24)

35 **Admin and MFA management UI refinements** IMPROVED 60

Adds Password Date (last password set time) as a viewable column on the Users page in the Admin Portal, splits WebAuthn enrollment into separate credential types (Security Keys and Device credentials), adds lock-screen approval/denial of JumpCloud Protect push notifications via long-press (iOS) or notification expansion (Android), and enables MFA by default for all new admin accounts logging in to the Admin Portal.

— Names each UI/security change individually Release Notes 2023 (2026-09-24)

THINNER COVERAGE BELOW

36 **Password Manager Desktop App storage capacity increase** IMPROVED 58

Raises item storage capacity to 13,000 items with autofill enabled, or up to 30,000 items without autofill; a single shared folder can hold up to 5,000 items.

— Exact numeric limits, no before value given What's New: Password Manager Desktop Ap…

37 **Shadow AI Dashboard and AI & SaaS Management rename** NEW 57

Introduces the Shadow AI Dashboard, a dedicated UI for visibility into generative AI app adoption, usage by department, and most active users, with the ability to set approvals or restrictions; JumpCloud SaaS Management is renamed AI & SaaS Management with expanded generative AI app coverage for higher discovery accuracy.

— Describes scope but no concrete config or metrics Release Notes 2025 (2026-09-24)

38 **RSS feed subscriptions for release notes** NEW 55

Adds RSS feed subscriptions for JumpCloud Help Center updates, product release notes, agent release notes, and bug fix reports via the Release Notes and Bug Fixes page, e.g. subscribing to only the Agent Release Notes feed.

Subscribe to only agent release notes to track JumpCloud agent changes without noise from other Help Center content.

> 📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click Subscribe to Agent Release Notes RSS Feed. 3. Copy the feed URL into your RSS reader of choice (e.g., Feedly, or a browser extension like RSS Feed Reader for Chrome).

– Exact page and steps to subscribe given `Subscribe to New Feature and Product Up…`

39 **Multi-Tenant Portal bulk user/device management** NEW 53

Adds ability for MSP admins to add new users and devices directly from the Multi-Tenant Portal, export users and devices lists, and apply bulk actions to users.

– Names actions but no export format or path given `Release Notes 2023 (2026-09-24)`

40 **Self-Service Account Provisioning at login window** NEW 52

Adds Self-Service Account Provisioning allowing new users to bind their account to JumpCloud-managed macOS and Windows devices directly from the login window, with a new macOS login window featuring wireless connectivity controls.

– Describes flow, no step-by-step or config given `Release Notes 2023 (2026-09-24)`

41 **ADI automatic AD security group creation** NEW 52

Automatically creates new security groups in AD from JumpCloud in the Root User Container specified for the integration, enabling complete security group management in AD from JumpCloud.

– Describes mechanism but no setup steps `product docs`

42 **ADI admin portal UX improvements** IMPROVED 52

Adds LDAP domain name formatting validation to the domain creation modal to prevent integration failures, shows the count of each agent type (sync and import) installed in the configuration sections of the Admin Portal Details tab, and saves the configuration selection on the Details tab so it persists after closing and reopening the domain instance configuration.

– Names the exact UI surfaces improved `product docs`

43 **Password Manager EU region support** NEW 51

Supports EU region selection before account activation — choose the region on the activation page (desktop app) or pairing page (browser extension).

— Names exact selection screens, limited further detail

What's New: Password Manager Desktop Ap…

44 **Password Manager launch** NEW 50

Adds JumpCloud Password Manager with end-to-end encrypted local sync, browser extension auto-fill, secure password sharing, and password generation.

— Describes capabilities but no config or limits Release Notes 2022 (2026-09-24)

45 **New SSO and access integrations: AWS Verified Access and Deel** NEW 50

Adds integration with AWS Verified Access using JumpCloud Go, SAML, or OIDC SSO as a VPN-less remote access solution, and adds a Deel SSO Connector integration with JumpCloud.

— Names both integrations, no setup detail Release Notes 2023 (2026-09-24)

46 **Apple Intelligence restriction MDM policies** NEW 48

Adds two new JumpCloud MDM policies to restrict Apple Intelligence features on iOS/iPadOS 18.1 and macOS 15.1 Sequoia.

— Names exact OS versions targeted, no policy detail Release Notes 2024 (2026-09-24)

47 **Password Manager ARM build for Apple Silicon** NEW 46

Provides a native ARM build for macOS devices with M-series silicon (M1 and above), with automatic migration from x64 builds.

— Names chip support and migration, no rollout detail

What's New: Password Manager Desktop Ap…

48 **AI Commands Builder reaches GA** NEW 45

JumpCloud AI Commands Builder reaches GA, generating and refining commands for Windows, macOS, and Linux devices from plain-English prompts within the Admin Portal.

— Explains behaviour but no example or config detail Release Notes 2025 (2026-09-24)

49 **Direct SSH access in Privileged Access Management** NEW 45

Adds 'Connect to a Resource with Direct Access using SSH' capability to Privileged Access Management, enabling SSH connections directly to managed resources.

— Names the capability, no config or limit detail product docs

50 **BYOD Account-Driven User Enrollment for iOS/iPadOS 18+** NEW

44

Adds general availability support for BYOD Account-Driven User Enrollment (ADUE) of iOS/iPadOS 18+ devices in MDM.

— States GA and OS scope, no enrollment steps Release Notes 2024 (2026-09-24)

51 **Expanded agent OS support: Fedora 37 and Windows 11 25H2**

IMPROVED

43

Adds JumpCloud Agent support for Linux Fedora 37, and adds Windows 11 25H2 support across JumpCloud's Windows management components.

— Names the two OS versions, no rollout detail Release Notes 2025 (2026-09-24)

Release Notes 2022 (2026-09-24)

52 **Windows Personalization policies** NEW

42

Adds new Windows Personalization policies: Taskbar Customization, Screen Saver Configuration, and Desktop and Lock Screen Settings.

— Names all three policies but no configuration detail Release Notes 2025 (2026-09-24)

53 **SaaS License Management reaches GA** NEW

41

Launches SaaS License Management (GA) with connectors for automated license tracking, usage optimization, and cost management.

— No named connectors or limits given Release Notes 2025 (2026-09-24)

54 **Legacy Password Manager CSV export for Vault migration**

IMPROVED

40

Enhances CSV export for the legacy Password Manager with new data fields to support migration into JumpCloud Vault.

— No field names or export steps given What's New: Password Manager Desktop Ap…

55 **TOTP MFA for LDAP applications** NEW

38

Adds TOTP as a supported MFA factor for LDAP applications, alongside JumpCloud Protect push notifications.

— Names the factor but no configuration steps Release Notes 2022 (2026-09-24)

56 **OIDC Single Sign-On support** NEW

38

Adds OpenID Connect (OIDC) Single Sign-On support so users can authenticate to OIDC-capable web applications via JumpCloud.

— No setup steps or scope detail given Release Notes 2022 (2026-09-24)

57 **Apple Business Manager service discovery for ADE** NEW

35

Adds Apple Business Manager as a new service discovery method for Account-driven Enrollment in Apple MDM.

— Single-line addition, no setup steps Release Notes 2025 (2026-09-24)

58 **Custom Admin Roles** NEW

33

Releases Custom Admin Roles for all direct organizations, enabling granular role tailoring based on specific responsibilities.

— No named permissions or roles listed Release Notes 2025 (2026-09-24)

59 **Service Accounts for APIs** NEW

33

Introduces Service Accounts for APIs that are not tied to individual Admin accounts, improving the security and reliability of integrations.

— Mechanism described only in general terms Release Notes 2025 (2026-09-24)

60 **IPA package support in Private Repository** NEW

32

Adds IPA file package support to the JumpCloud Private Repository for deploying iOS and iPadOS apps.

— Single-line addition with no workflow detail Release Notes 2025 (2026-09-24)

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! Staged users in JumpCloud are no longer provisioned to AD — users that were previously created in AD from the staged state will no longer be created until their state is changed to active.

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; replace with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups API `memberQuery` filter operators `$le` and `$ge` are deprecated and replaced by `$lte` and `$gte`; existing API scripts using `$le` / `$ge` must be updated.

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

JumpCloud: Adds a unified Patch Management dashboard in the Admin Portal with `Overview`, `Windows`, and `Apple` tabs for cross-platform patch compliance, severity filtering, and bulk remediation.

└──▷ **HOW TO FIND IT**

Surface Windows KB patch gaps across your fleet — filter the dashboard to Critical severity, then bulk-install missing patches immediately instead of waiting for the next policy window.

> 📍 In the Admin Portal, go to Patch Management › Windows, set the Global severity filter to `Critical`, select non-compliant devices, and choose `Install missing critical patches` › `Install immediately`.

Push an updated application binary to already-enrolled devices without reassigning device groups or recreating the app entry in the Private Repository.

> 📍 In the Admin Portal, go to the application's Details tab in the Private Repository, click `Update Version`, upload the new file via the `Upload File` flow, enter an install success exit code (for EXE), and confirm deployment.

Ground AI Assistant in org-specific runbooks and SOPs so it can answer questions about internal processes, with answers attributed to the source document.

> 📍 In the Admin Portal, navigate to AI Assistant › Knowledge Base, click `Upload`, and add your internal PDF, TXT, DOCX, or MD files (up to 10 MB each, 25 files per org).

Adds Windows-specific patch dashboard with KB-centric controls: search/filter KBs, review CVE details, approve awaiting updates, retry failed installations, and export the KB list as CSV or JSON.

Adds Apple Compliance Dashboard tracking macOS, iOS, and iPadOS OS rollout progress against active update policies, with filtering by platform, release date, severity, and device status.

Adds `Update Version` action on an existing application's `Details` tab in the Private Repository, enabling binary updates without re-creating the application or reassigning devices.

Adds new Android EMM policies for Display Settings, Autofill (Android 8.0+), Work Account (Google enterprise account enforcement), and Default App and Intent Handling.

`Release Notes 2022 (2026-09-24)` `BLOG ↗` `STABLE`

JumpCloud 2022 adds passwordless RADIUS auth, TOTP for LDAP, Remote Assist, Password Manager, and new MSP command templates.

└──▷ **HOW TO FIND IT**

Block device access via a Conditional Access Policy when disk encryption is not enabled, covering FileVault (macOS), BitLocker (Windows), and root disk encryption (Linux).

📍 In the Admin Portal, go to Security › Conditional Access › New Policy, add a Device Condition, select `Disk Encryption`, set the condition to `disabled`, and set the policy action to `Deny`.

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 certificate (passwordless) based authentication to JumpCloud RADIUS, supporting certificates from any X.509 compliant certificate authority.

Adds TOTP as a supported MFA factor for LDAP applications, alongside JumpCloud Protect push notifications.

Adds a disk-encryption device condition to Conditional Access Policies, allowing or denying access based on FileVault (macOS), system drive encryption (Windows), or root disk encryption (Linux).

Adds OpenID Connect (OIDC) Single Sign-On support so users can authenticate to OIDC-capable web applications via JumpCloud.

+17 more

Release Notes 2023 (2026-09-24) BLOG ↗ STABLE

JumpCloud 2023 adds passwordless Go auth, Android zero-touch enrollment, Remote Assist background tools, dynamic groups, and Push Bombing mitigation.

└──▷ **HOW TO FIND IT**

Tighten admin-portal access by verifying that MFA is enforced for new admin accounts and that Push Bombing mitigation is active — check or adjust the concurrent-attempt limit in MFA Configurations.

📍 In the Admin Portal, go to Security › MFA Configurations, confirm Push Bombing mitigation is enabled, and set the maximum concurrent push attempts appropriate for your org.

Automate device group membership so newly enrolled corporate laptops running a specific OS version are added to the right policy group without manual review.

📍 In the Admin Portal, go to Device Groups › select or create a group › enable Dynamic › add attribute rules (e.g., Vendor = `Apple`, OS Major Version = `14`) › save to activate automatic membership updates.

Remotely troubleshoot an end-user machine during an incident without interrupting the user's active session, using Background Tools command line.

> 📍 In the Admin Portal, go to Devices › select the target device › Background Tools › Remote Command Line, then run your diagnostic commands without launching a visible Remote Assist session.

Adds `JumpCloud Go` browser extension enabling hardware-protected, phishing-resistant passwordless login to JumpCloud-protected web resources on managed macOS and Windows devices via Apple Touch ID and Windows Hello.

Adds `JumpCloud Go` as the default MFA method for conditional access policies on SSO apps, with user and device verification on every application session.

Adds Android zero-touch enrollment support for fully managed, dedicated, and company-owned Android devices with a work profile.

Adds Remote Assist Background Tools enabling admins to use remote command line and remote file transfer to access end-user devices without disrupting active sessions.

Adds Self-Service Account Provisioning allowing new users to bind their account to JumpCloud-managed macOS and Windows devices directly from the login window, with a new macOS login window featuring wireless connectivity controls.

+21 more

Release Notes 2024 (2026-09-24) BLOG ↗ STABLE

JumpCloud 2024 adds Device Monitoring & Alerting, Windows App Catalog, SaaS Management, CSP-based MDM policies, and multi-tenant MFA enforcement.

└──▷ **TRY IT**

Filter users whose department is `Finance | Accounting` or `IT` using the newly escapable pipe character in API filter search terms.

```
$ curl -s -X GET 'https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT' \
  -H 'x-api-key: <your-api-key>' \
  -H 'Content-Type: application/json'
```

Migrate an automation that reactivates a user account from the deprecated GET endpoint to the new POST endpoint.

```
$ curl -s -X POST
'https://console.jumpcloud.com/systemusers/reactivate/{id}' \
  -H 'x-api-key: <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d '{"activation_email": "user@example.com"}'
```

Enforce MFA for all MSP admins across a multi-tenant portal from a single billing admin account.

📍 In the Admin Portal, go to Account › Settings and enable the `Global MFA Requirement` toggle.

Adds `GET /systemusers?filter=department:$in:Finance | Accounting|IT` support: the pipe `|` character can now be escaped with `\` inside `in` and `nin` API filter operator search terms.

Adds six new Windows MDM policies via Configuration Service Provider (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.

Adds `$lte` and `$gte` filter operators for Dynamic Groups API `memberQuery` filters, replacing the deprecated `$le` and `$ge` operators.

Adds `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` endpoint to replace the deprecated `GET api/systemusers/reactivate/{id}` request (which now returns 404).

`+16 more`

└─▷ **BREAKING ON UPGRADE**

- ! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; replace with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.
- ! Dynamic Groups API `memberQuery` filter operators `$le` and `$ge` are deprecated and replaced by `$lte` and `$gte` ; existing API scripts using `$le` / `$ge` must be updated.

Release Notes 2025 (2026-09-24) BLOG ↗ STABLE

JumpCloud ships Access Requests with webhook/JIRA integration, MCP Server GA, Shadow AI Dashboard, forced re-auth SAML/OIDC, and a wave of new management features.

└──▷ **HOW TO FIND IT**

Integrate Access Requests with JIRA so a ticket is automatically created whenever a user submits an access request, enabling your existing approval workflow.

> 📍 In the JumpCloud Admin Portal, go to Access Requests › Settings, create an approval flow, and configure a webhook pointing to your JIRA project's issue creation endpoint to receive real-time notifications on new requests.

Restrict a group of remote users in regions without Push support to TOTP-only MFA while keeping other groups on JumpCloud Protect Push.

> 📍 In the JumpCloud Admin Portal, go to MFA › JumpCloud Protect › Unified MFA settings, and add the target user group to the Push exclusion list.

Adds `sso_auth` event to logs to capture forced re-authentication events for compliance and security auditing.

Adds `ForceAuthn='true'` (SAML) and `prompt=login` (OIDC) support so that service providers can trigger step-up authentication without ending the user's existing session.

Includes Authentication Methods References (AMR) claims in both SAML assertions and OIDC ID tokens to surface proof of MFA factors used at the point of access.

Introduces JumpCloud Access Requests, enabling configurable approval flows that users submit from the JumpCloud User Portal and admins manage in the console.

Launches the JumpCloud MCP Server (GA), allowing compatible AI tools and code editors to interact with organizational data and administrative capabilities via the Model Context Protocol.

+22 more

Release Notes 2026 (2026-09-24) BLOG ↗ STABLE

JumpCloud adds Password Vault (admin, user, mobile, browser), Agent Identities, MCP visibility, and expanded Android/Windows policies in 2026-09.

└──▷ **HOW TO FIND IT**

Surface AI agents deployed in your environment and connect a newly discovered MCP server to AI Gateway for centralized access control.

> 📍 In the Admin Portal, go to Access › AI & SaaS Management › Applications. Check the MCP Connection column for any app showing Discovered, click the status, review

the listed MCP servers, then click Add MCP Server to register it in AI Gateway.

Register a new AI agent identity and issue it OAuth credentials scoped only to the SSO applications it needs, keeping agent access auditable and separate from human accounts.

📍 In the Admin Portal, go to Identity Management › Agents and click + Add. Fill in the agent name, owner, and company email on the Details tab. On the Authentication tab, generate OAuth client credentials and set the client authentication type to `Client secret post` if the target platform (e.g. Amazon Bedrock AgentCore) requires credentials in the request body. Then go to Identity Management › Agent Groups, create a group, add the agent, and assign the required SSO applications.

Audit Windows patch compliance across all managed devices without any manual data collection — download the full result set for offline analysis or ticketing.

📍 In the Admin Portal, navigate to JumpCloud Reports › Device Patch Report (Windows). Preview up to 500 rows to spot devices with `Failed` or `Pending` KB status, then click Download to export the complete organization result set as CSV.

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 `Identity Management > Agents` section (Public Preview) to register and govern non-human AI agent identities with OAuth-based credentials, group-based application access, and lifecycle management independent of human accounts.

Adds Password Vault to the JumpCloud Admin Portal, enabling admins to enable the vault, assign access to users and user groups, create and share credentials and websites with `Connect`, `View`, `Edit`, and `Manage` permissions, and view audit logs in Directory Insights.

Adds Password Vault to the JumpCloud User Portal (left navigation), letting users manage and share credentials and websites, bulk-import credentials via CSV or Excel, and organize items in personal and shared folders.

Adds Password Vault to the JumpCloud Protect mobile app so users can view secrets, autofill credentials into mobile apps, and launch websites with autofill — authentication retrieves credentials from JumpCloud while TOTP codes remain stored locally.

`+3 more`

`Subscribe to New Feature and Product Update Announcements (2026-09-24)` `BLOG ↗`

`STABLE`

JumpCloud adds RSS feed subscriptions for release notes, agent updates, and bug fix announcements.

└──▷ **HOW TO FIND IT**

Subscribe to only agent release notes to track JumpCloud agent changes without noise from other Help Center content.

📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click `Subscribe to Agent Release Notes RSS Feed`. 3. Copy the feed URL into your RSS reader of choice (e.g., Feedly, or a browser extension like RSS Feed Reader for Chrome).

`What's New: Password Manager Desktop App (2026-09-24)` **BLOG ↗** **STABLE**

JumpCloud Password Manager Desktop App adds EU region support, ARM builds for Apple Silicon, auto-lock, and raises item storage to 30,000.

└──▷ **HOW TO FIND IT**

Lock down unattended workstations by enabling auto-lock so the Password Manager vault locks whenever the device itself locks or sleeps.

📍 In the Password Manager desktop app, go to Settings › Auto Lock and enable the `Auto Lock` toggle to lock the app when the device is locked.

Supports EU region selection before account activation — choose your region on the activation page (desktop app) or pairing page (browser extension).

Introduces auto-lock schedule with a default of 15 minutes of inactivity and a configurable maximum of 1,440 minutes (24 hours).

Raises item storage capacity to 13,000 items with autofill enabled, or up to 30,000 items without autofill; a single shared folder can hold up to 5,000 items.

Adds `Move All` button in shared folders to move all items from a folder in one action.

Paginates the items list at 100 items per page for improved navigation at scale.

`+14 more`

└──▷ **BREAKING ON UPGRADE**

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ

New Availability: The SailPoint Identity Graph is Coming to Your Environment (2026-09-23)

BLOG ↗ RANK

SailPoint is an identity governance and access management platform that automates user provisioning, access reviews, and compliance across enterprise systems.

> SailPoint's biggest window addition is the Identity Graph, now rolling out to all Business+ customers for interactive identity relationship and blast-radius visualization, alongside a new workflow version History tab and connectors for over twenty external systems.

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

01 History tab for NERM Classic workflow versioning NEW

85

NERM Classic workflow configuration gains a History tab showing a paginated, newest-first table of version-lifecycle and step/configuration-change events — including when, who, event type, step label, and version identifier — without requiring Admin Activity or the audit API. The tab supports search by step labels, step IDs, and version IDs with quick filters to narrow the event list, and enables version comparison from within the tab, letting admins follow changes that spawned a version and open a diff, including after a fork.

Review the full change history of a NERM Classic workflow after a suspected misconfiguration, filtering by version ID to isolate the offending change.

📍 1. Open the target workflow in NERM Classic Admin. 2. Select the History tab. 3. Enter the version ID in the search box to filter events. 4. Select the matching Version Created or Version Forked row to open event details. 5. Click through to the version comparison view to see exactly what changed.

— Exact UI steps and searchable fields make this immediately usable

Enhancement: Workflow Versioning — Hist…

02 Identity Graph across MySailPoint and context menus NEW

70

SailPoint Identity Graph rolls out to all Business+ customers, accessible directly from the MySailPoint homepage and from context menus across the product by clicking the graph icon on any identity, entitlement, or access item. It visualizes the full identity estate — humans, non-human identities, and agents — in a single interactive graph to surface overpermissioned identities and complex access pathways, identifies the blast

radius of a compromised identity using context across SailPoint's portfolio, and enables direct remediation actions from within the visualization without leaving it.

Navigate to Identity Graph for a specific identity or entitlement directly from any list or menu in the product.

> 📍 In any identity, entitlement, or access list across the SailPoint product, click the Identity Graph icon in the top right of the item row. Identity Graph loads scoped to that item, showing its full access context and relationships.

— Describes mechanism and access path but no API or config surface

New Availability: The SailPoint Identit…

New Availability: The SailPoint Identity Graph is Coming to Your Environment (2026-09-

BLOG ↗ LATEST STABLE

SailPoint Identity Graph rolls out to all Business+ customers, enabling interactive visualization of identity relationships and blast radius analysis.

└──▷ **HOW TO FIND IT**

Navigate to Identity Graph for a specific identity or entitlement directly from any list or menu in the product.

> 📍 In any identity, entitlement, or access list across the SailPoint product, click the Identity Graph icon in the top right of the item row. Identity Graph loads scoped to that item, showing its full access context and relationships.

Accesses Identity Graph directly from the MySailPoint homepage and from context menus across the product by clicking the graph icon on any identity, entitlement, or access item.

Visualizes the full identity estate — humans, non-human identities, and agents — in a single interactive graph to surface overpermissioned identities and complex access pathways.

Identifies the blast radius of a compromised identity immediately, leveraging context across SailPoint's portfolio.

Enables direct remediation actions from within Identity Graph to reduce risk without leaving the visualization.

Enhancement: Workflow Versioning — History tab (2026-09-23) BLOG ↗ STABLE

NERM Classic workflows gain a History tab for paginated, searchable version and change event review.

└──▷ **HOW TO FIND IT**

Review the full change history of a NERM Classic workflow after a suspected misconfiguration, filtering by version ID to isolate the offending change.

📍 1. Open the target workflow in NERM Classic Admin. 2. Select the `History` tab. 3. Enter the version ID in the search box to filter events. 4. Select the matching `Version Created` or `Version Forked` row to open event details. 5. Click through to the version comparison view to see exactly what changed.

Adds a History tab to NERM Classic workflow configuration, showing a paginated newest-first table of version-lifecycle and step/configuration-change events — including when, who, event type, step label, and version identifier — without requiring Admin Activity or the audit API.

Supports search by step labels, step IDs, and version IDs within the History tab, with quick filters to narrow the event list.

Enables version comparison from within the History tab, letting admins follow changes that spawned a version and open a diff, including after a fork.

WAS THIS USEFUL? 👍 👎

Aembit

Feedback

changelog · Sep 24, 2026

CHANGELOG ↗

RANK

Aembit is IAM for agentic AI and workloads. Enable secretless access and real-time policy enforcement across any environment to protect data and accelerate AI adoption.

> Aembit's biggest addition this window is MCP Tool Access Control for the MCP Identity Gateway, paired with new X.509-SVID certificate delivery through the Secrets Operator, SAML-bound dynamic claims, multi-cluster Kubernetes attestation, and AWS STS Credential Provider selection extended across GitHub Actions and the Edge SDKs.

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

01 **X.509-SVID delivery via Secrets Operator**

BREAKING

95

New `credentialType: X509Svid` lets the Secrets Operator issue mTLS certificates to a workload without the private key ever leaving the cluster, handling CSR generation and automatic renewal. Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret — X509Svid writes a `kubernetes.io/tls` Secret while all other types write a generic Secret.

Issue mTLS certificates to a workload without the private key ever leaving the cluster — Secrets Operator handles CSR generation and automatic renewal.

yaml

```yaml
credentialType: X509Svid
```

— Exact config key, secret type and migration behavior spelled out

changelog-20260924-58a29786

02 **AWS STS Credential Provider selection in GitHub Actions and Edge SDKs**

NEW

95

Adds `aws-access-key-id` input to the Aembit GitHub Action (minimum version 1.3.0) so a workflow on a GitHub-hosted runner can select among multiple AWS STS Federation Credential Providers in a single Access Policy by Access Key ID selector. Extends the Edge API credentials endpoint and Edge SDKs for TypeScript and Python (minimum Edge SDK 1.34.0) to pass an AWS STS Credential Provider selector as connection metadata alongside existing JWT Credential Provider selectors.

Let a GitHub-hosted Actions runner assume a specific IAM role from an Access Policy that holds multiple AWS STS Federation Credential Providers, identified by Access Key ID selector.

```
$ - uses: aembit/aembit-github-action@v1
  with:
    aws-access-key-id: AKIA<SELECTOR_FOR_TARGET_ROLE>
```

— Exact flag, minimum versions and runnable workflow snippet changelog-20260924-58a29786

03 **MCP Tool Access Control in MCP Identity Gateway** NEW

85

MCP Identity Gateway (version 1.34.5794) adds Tool Visibility, which controls which MCP tools an AI agent discovers, and Tool Invocation, which controls which tools it can call; both accept exact tool names or case-sensitive glob patterns. New `application.mcp.contentSecurity` fields on MCP workload authorization events record which control decided a tool's fate, which rules matched, which tools a listing removed, and how many it returned.

— Names exact fields and gateway version but no config path shown

changelog-20260924-58a29786

04 **SAML-bound dynamic claims in OAuth Credential Providers** NEW

75

A credential can now be bound to the authenticated SAML user's identity in an OAuth 2.0 Authorization Code Credential Provider using `${saml.response.subject.nameId}`, so each user gets a token scoped to their own account.

Bind a credential to the authenticated SAML user's identity in an OAuth 2.0 Authorization Code Credential Provider, so each user gets a token scoped to their own account.

yaml

```
${saml.response.subject.nameId}
```

— Exact binding syntax given, no wider config context changelog-20260924-58a29786

THINNER COVERAGE BELOW

05 **Windows certificate chain validation in Agent Controller**

IMPROVED

45

Agent Controller now installs intermediate CA certificates from its TLS certificate chain into the Windows LocalMachine certificate store, so Windows TLS clients can validate the full certificate chain.

— Mechanism explained, no manual step needed or given changelog-20260924-58a29786

06 **SSO user profile population from OIDC claims** IMPROVED

45

When Aembit creates a user through single sign-on, it now populates the user's first and last names from OIDC profile claims `given_name` and `family_name`, falling back to `name` if those are absent.

— Exact claim fields named, no admin-facing action changelog-20260924-58a29786

07 Kerberos Trust Provider attestation for Agent Proxy on EC2

NEW

35

Agent Controller (version 1.34.3725) adds support for Kerberos Trust Provider attestation requests from Agent Proxy when running on AWS EC2 instances.

— Version and scope named, no configuration detail changelog-20260924-58a29786

08 Multi-cluster Kubernetes attestation NEW

10

Aembit adds multi-cluster Kubernetes attestation, per the release summary; no further mechanism or configuration detail is provided in the changelog entries.

— Only named in headline, no supporting detail given changelog-20260924-58a29786

└──▷ **BREAKING ON UPGRADE**

! Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret (X509Svid writes a `kubernetes.io/tls` Secret; all other types write a generic Secret).

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Aembit adds SAML-bound dynamic claims, X.509-SVID delivery, multi-cluster Kubernetes attestation, MCP Tool Access Control, and AWS STS provider selection in GitHub Actions.

└──▷ **USE IT**

Issue mTLS certificates to a workload without the private key ever leaving the cluster — Secrets Operator handles CSR generation and automatic renewal.

yaml

```
credentialType: X509Svid
```

Bind a credential to the authenticated SAML user's identity in an OAuth 2.0 Authorization Code Credential Provider, so each user gets a token scoped to their own account.

yaml

```
${saml.response.subject.nameId}
```

Let a GitHub-hosted Actions runner assume a specific IAM role from an Access Policy that holds multiple AWS STS Federation Credential Providers, identified by Access Key ID selector.

```
$ - uses: aembit/aembit-github-action@v1
  with:
    aws-access-key-id: AKIA<SELECTOR_FOR_TARGET_ROLE>
```

Adds `aws-access-key-id` input to the Aembit GitHub Action (minimum version 1.3.0) so a workflow on a GitHub-hosted runner can select among multiple AWS STS Federation Credential Providers in a single Access Policy by Access Key ID selector.

Extends the Edge API credentials endpoint and Edge SDKs for TypeScript and Python (minimum Edge SDK 1.34.0) to pass an AWS STS Credential Provider selector as connection metadata alongside existing JWT Credential Provider selectors.

Adds MCP Tool Access Control enforcement to the MCP Identity Gateway (version 1.34.5794): Tool Visibility controls which MCP tools an AI agent discovers, Tool Invocation controls which it can call, and both accept exact tool names or case-sensitive glob patterns.

Adds `application.mcp.contentSecurity` fields to MCP workload authorization events, recording which control decided a tool's fate, which rules matched, which tools a listing removed, and how many it returned.

Adds Agent Controller support for Kerberos Trust Provider attestation requests from Agent Proxy when running on AWS EC2 instances (version 1.34.3725).

+2 more

└─▷ **BREAKING ON UPGRADE**

! Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret (X509Svid writes a `kubernetes.io/tls` Secret; all other types write a generic Secret).

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ changelog · Sep 24, 2026

CHANGELOG ↗ RANK

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across infrastructure.

> DigiCert CertCentral will shift default TLS issuance to new G5 root hierarchies, tightens domain validation to meet CA/Browser Forum MPIC requirements, and adds a persistent DNS TXT method for automatic recurring domain revalidation.

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

01 Default TLS issuance moves to G5 root hierarchies

BREAKING

73

As of October 15, 2026, new orders, renewals, reissues, and duplicates for DV, OV, and EV certificates across DigiCert, GeoTrust, Thawte, RapidSSL, and Encryption Everywhere brands will be issued by default from the `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` hierarchies. Customers who hard-code or pin root/intermediate CA certificates, or manage their own trust stores, must add and trust the new G5 root and intermediate CA certificates or face service disruptions.

— Names exact root/hierarchy identifiers, brands, and effective date.

changelog-20260924-6c9a2c52

02 Stricter MPIC domain validation enforcement

BREAKING

70

Effective December 1/15, 2026, CertCentral enforces six remote MPIC network perspectives across at least two RIR regions for domain control verification and CAA checks, requiring four corroborations per domain query, in compliance with CA/Browser Forum MPIC requirements. Environments that restrict DNS queries or website traffic by IP address, region, query volume, or rate must verify their configurations support traffic from six remote perspectives or certificate issuance may fail.

— Specifies perspective count, region requirement, corroboration threshold, and date.

changelog-20260924-6c9a2c52

└──▷ **ALSO FROM THESE RELEASES**

└──▷ **BREAKING ON UPGRADE**

! TLS certificates (new orders, renewals, reissues, duplicates) will be issued from `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` by default as of October 15, 2026 — customers who hard-code or pin root/intermediate CA certificates, or manage their own trust stores, must add and trust the G5 root and intermediate CA certificates or face service disruptions.

! MPIC enforcement as of December 1, 2026 increases to six remote network perspectives requiring four corroborations — environments that restrict DNS queries or website traffic by IP address, region, query volume, or rate must verify their configurations support MPIC traffic from six remote perspectives or certificate issuance may fail.

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

CertCentral adds Persistent DNS TXT automatic domain revalidation and shifts TLS issuance to G5 RSA/ECC roots by default

└──▷ **HOW TO FIND IT**

Set up Persistent DNS TXT automatic domain revalidation so domains stay continuously validated without manual intervention — useful for teams issuing OV/EV TLS certificates on recurring schedules.

📍 In CertCentral, go to Domains, select the domain you want to keep continuously validated, choose Persistent DNS TXT Record as the DCV method, complete the initial validation, then leave the Persistent DNS TXT record in place with the correct persistent URI value (Account or Unique). CertCentral will revalidate the domain automatically every five days.

Shifts default TLS certificate issuance (new orders, renewals, reissues, and duplicates) to `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` hierarchies for all DV, OV, and EV certificates across DigiCert, GeoTrust, Thawte, RapidSSL, and Encryption Everywhere brands.

Enforces six remote MPIC network perspectives across at least two RIR regions for domain control verification and CAA checks, requiring four corroborations per domain query, in compliance with the CA/Browser Forum MPIC requirements effective December 15, 2026.

└──▷ **BREAKING ON UPGRADE**

! TLS certificates (new orders, renewals, reissues, duplicates) will be issued from `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` by default as of October 15, 2026 — customers who hard-code or pin root/intermediate CA certificates, or manage their own trust stores, must add and trust the G5 root and intermediate CA certificates or face service disruptions.

! MPIC enforcement as of December 1, 2026 increases to six remote network perspectives requiring four corroborations — environments that restrict DNS queries or website traffic by IP address, region, query volume, or rate must verify

their configurations support MPIC traffic from six remote perspectives or certificate issuance may fail.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Varonis

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

The leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis's latest changelog centers on a new Data Lifecycle Management capability for automatically triaging redundant and obsolete data, alongside new integration and deployment options including HashiCorp Vault-backed collector credentials, Snowflake connectivity over AWS PrivateLink, Terraform-based collector deployment, Confluence history scanning, and Jira permissions analysis.

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

01 HashiCorp Vault secret manager for collector credentials NEW

80

Route all Private Cloud Data Collector credentials through your existing HashiCorp Vault instead of storing them in Varonis, to centralize access control and meet compliance requirements. Configure via Settings > Data Collectors, select your Private Cloud Data Collector, choose Secret Manager, select HashiCorp Vault, then enter your Vault endpoint and authentication details.

Route all Private Cloud Data Collector credentials through your existing HashiCorp Vault instead of storing them in Varonis, to centralize access control and meet compliance requirements.

📍 In the Varonis console, go to Settings › Data Collectors, select your Private Cloud Data Collector, choose Secret Manager, and select HashiCorp Vault. Enter your Vault endpoint and authentication details to complete the configuration.

— Exact UI navigation and config steps given changelog-20260924-213595ed

02 Data Lifecycle Management for ROT data NEW

60

Introduces Data Lifecycle Management (DLM) to automatically identify, prioritize, and quarantine redundant, obsolete, and trivial (ROT) data across connected data sources using classification, permissions mapping, and activity analysis.

— Explains mechanism but gives no UI path or config to act on changelog-20260924-213595ed

THINNER COVERAGE BELOW

03 Terraform-based collector deployment NEW

25

Adds Terraform-based deployment for Varonis data collectors, per the release summary.

— Only named in summary; no mechanism or steps given changelog-20260924-213595ed

04 **Jira permissions analysis** NEW 25

Adds permissions analysis for Jira, per the release summary.

— Only named in summary; no further detail provided changelog-20260924-213595ed

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Varonis adds Data Lifecycle Management, Terraform-based collector deployment, Confluence history scanning, Jira permissions analysis, and more.

└──▷ **HOW TO FIND IT**

Route all Private Cloud Data Collector credentials through your existing HashiCorp Vault instead of storing them in Varonis, to centralize access control and meet compliance requirements.

📍 In the Varonis console, go to Settings › Data Collectors, select your Private Cloud Data Collector, choose Secret Manager, and select HashiCorp Vault. Enter your Vault endpoint and authentication details to complete the configuration.

Connect Varonis to a Snowflake environment routed over AWS PrivateLink without leaving the Varonis UI, validating connectivity before the integration goes live.

📍 In the Varonis console, go to Integrations › Add Integration › Snowflake, select AWS PrivateLink as the connection type, allow the wizard to generate the endpoint URL, and click Validate before saving.

Introduces Data Lifecycle Management (DLM) to automatically identify, prioritize, and quarantine redundant, obsolete, and trivial (ROT) data across connected data sources using classification, permissions mapping, and activity analysis.

WAS THIS USEFUL? 👍 👎

Ubiq Security

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

RANK

Learn how Ubiq discovers sensitive data, maps identity-to-data access, and enforces protection policies across applications, databases, APIs, and AI.

> Ubiq shipped new Airflow operators for column-level encryption, searchable structured encryption and identity-provider authentication in the JavaScript SDK, a key-number lookup API in the Go SDK, a Grafana Loki SIEM integration, and broadened character-set support for date/datetime datasets.

└──▷ **WHAT SHIPPED · 7 FEATURES** most completely described first ? what's the number?

01 Airflow operators for column-level encryption NEW

95

New `UbiqEncryptOperator` and `UbiqDecryptOperator` Airflow operators accept a `credentials` parameter (Ubiq credentials object), a `data_task` parameter (task ID of the upstream data task), and a `dataset_names` list mapping columns to Ubiq dataset names, enabling column-level encryption on tabular data such as BigQuery results. Install via `pip install ubiq_security`; operators are loaded automatically when placed under `plugins/operators/` in the Airflow home directory.

Encrypt sensitive columns (SSN, Birthdate) pulled from BigQuery before storing or forwarding them, then decrypt in the same DAG to verify round-trip correctness.

```python
import ubiq_security as ubiq
from operators.ubiq_encrypt_operator import UbiqEncryptOperator
from operators.ubiq_decrypt_operator import UbiqDecryptOperator

creds = ubiq.configCredentials(config_file='ubiq-credentials',
profile='default')

get_data = BigQueryGetDataOperator(
    task_id='get_data',
    dataset_id=DATASET_NAME,
    table_id=TABLE_NAME,
    max_results=10,
    selected_fields='name,email,ssn,birthday',
)

enc_task = UbiqEncryptOperator(
    task_id='enc_task',
    credentials=creds,
    data_task=get_data.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate'],
)

dec_task = UbiqDecryptOperator(
    task_id='dec_task',
    credentials=creds,
    data_task=enc_task.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate'],
)

get_data >> enc_task >> dec_task
```

— Names exact operators, parameters, install command and directory. product docs

02

Encrypted search support for structured encryption in JavaScript SDK

NEW

85

The JavaScript library's structured encryption builder adds `encryptNumberForSearchAsync` alongside `EncryptNumberAsync`, producing all ciphertext variants needed to run encrypted search queries against a pre-configured dataset (e.g. an `int32` field) while keeping the value searchable in the database.

Encrypt an integer field (e.g. a customer ID) using a pre-configured 32-bit dataset so it remains searchable in the database.

javascript

```
const credentials =
ubiq.UbiqFactory.readCredentialsFromFile('./credentials',
'default');
const configuration = ubiq.UbiqFactory.defaultConfiguration();
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

const tweakFF1 = [];
const int32 = await
ubiqEncryptDecrypt.EncryptNumberAsync('int32', 9, tweakFF1);
const searchValues = await
ubiqEncryptDecrypt.encryptNumberForSearchAsync('int32', 9,
tweakFF1);
// searchValues contains all ciphertexts needed for an encrypted
search
```

— Exact method names and a runnable example provided. product docs

03 **Key number lookup without decryption in Go SDK** NEW

80

The Go SDK exposes ubiq.GetKeyNumber(ciphertext) to determine which key protects a given ciphertext without performing a full decrypt, intended for key-rotation audits.

Inspect which key number protects a ciphertext in Go — useful for key-rotation audits — without triggering a full decrypt.

go

```
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Protected by key number: %d\n", keyNum)
```

Inspect the key number behind an existing ciphertext in Go without decrypting it — useful for key rotation audits.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Println('Key number:', keyNum)
```

— Names exact function and use case with runnable code. changelog-20260924-917e6c4d changelog-20260924-9284cc01

04 Custom metadata tagging on encryption usage records NEW

75

`addReportingUserDefinedMetadata` lets callers attach arbitrary JSON context (e.g. `{"pipeline": "ingest", "team": "data-eng"}`) to an encryption/decryption call so the audit log can be correlated with external incident or pipeline systems.

Tag usage records with incident or pipeline context so the audit log can be correlated with external systems.

```javascript
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

ubiqEncryptDecrypt.addReportingUserDefinedMetadata('{"pipeline":
"ingest", "team": "data-eng"}');
const ct = await ubiqEncryptDecrypt.EncryptAsync('token64',
'abc', []);
```

— Exact method and example payload given. product docs

05 Identity provider authentication in JavaScript SDK NEW

70

Adds ubiq.UbiqFactory.createCredentialsWithIdp(username, password, host) so the JS SDK can authenticate via an identity provider instead of static API keys, for SSO-managed deployments.

Authenticate using an identity provider instead of static API keys when deploying in an SSO-managed environment.

```javascript
const credentials =
ubiq.UbiqFactory.createCredentialsWithIdp(<username>, <password>,
<host>);
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)

.withConfigurationObject(ubiq.UbiqFactory.defaultConfiguration())
  .buildStructuredAsync();
```

— Exact factory method named with runnable usage. product docs

06 **Grafana Loki SIEM integration** NEW 60

Ubiq added a unified Integrations page (Settings > Integrations) where Grafana Loki can be enabled as a SIEM destination for Ubiq activity events, by supplying a Loki endpoint and credentials.

Enable Grafana Loki as a SIEM destination for Ubiq activity events — navigate the console to reach the new unified Integrations page.

> 📍 In the Ubiq console, go to Settings › Integrations, select Grafana Loki under SIEM, and follow the prompts to supply your Loki endpoint and credentials.

— Names exact UI path but not an API or config key. changelog-20260924-917e6c4d

THINNER COVERAGE BELOW

07 **Character set support for date and datetime datasets** IMPROVED 40

The `date` and `datetime` format-preserving encryption datasets now support input character sets other than `0-9`, widening the range of tabular data that can be encrypted in place.

— Only a high-level description, no configuration example given.

changelog-20260924-917e6c4d

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Ubiq Security: `date` and `datetime` datasets now work with input character sets other than `0-9`.

└──▷ USE IT

Inspect which key number protects a ciphertext in Go — useful for key-rotation audits — without triggering a full decrypt.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Protected by key number: %d\n", keyNum)
```

Enable Grafana Loki as a SIEM destination for Ubiq activity events — navigate the console to reach the new unified Integrations page.

> 📍 In the Ubiq console, go to Settings › Integrations, select `Grafana Loki` under SIEM, and follow the prompts to supply your Loki endpoint and credentials.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale's mobile app now supports pinging devices and shows guidance when a tailnet is empty, while the admin console gains detailed posture-assertion views and webhook event subscriptions were introduced.

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

01 Posture assertion detail view in admin console NEW

70

In the admin console, open Machines > select a device > Device posture to view each posture assertion, the machine's reported value, and the count of policy rules requiring that posture — useful for checking compliance before pushing a policy change.

Review which posture assertions a specific machine is passing or failing before pushing a policy change.

> 📍 In the admin console, open Machines › select a device › Device posture — view each assertion, the machine's reported value, and the count of policy rules requiring that posture.

— Exact UI navigation path given, but no API/config surface changelog-20260924-60189015

THINNER COVERAGE BELOW

02 Device ping and empty-state banner in mobile app NEW

35

The mobile app can now ping devices in your tailnet directly, and it shows a helpful banner when there are no other devices in the tailnet.

— Names two mobile app additions but no mechanism detail changelog-20260924-60189015

03 Webhook event subscriptions NEW

10

Tailscale added webhook event subscriptions, per the release summary, though no further detail on configuration or supported events is given.

— Mentioned only in summary, no mechanism or usage detail changelog-20260924-60189015

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

> Tailscale adds posture-assertion detail views, webhook event subscriptions, and per-device ping from the mobile app.

└──▷ HOW TO FIND IT

Review which posture assertions a specific machine is passing or failing before pushing a policy change.

> 📍 In the admin console, open Machines › select a device › Device posture — view each assertion, the machine's reported value, and the count of policy rules requiring that posture.

Enables pinging devices in your tailnet directly from the mobile app.

Shows a helpful banner when there are no other devices in the tailnet.

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone consolidated gateway management into a single `firezone-gateway` binary and expanded device attestation to recognize an additional SAN claim alias.

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

01 Gateway CLI consolidated into firezone-gateway

BREAKING

85

The `authenticate`, `enable-service`, and `disable-service` subcommands moved from the separate `firezone` CLI into `firezone-gateway` itself, consolidating gateway management into a single binary. Any scripts or service definitions calling `firezone authenticate`, `firezone enable-service`, or `firezone disable-service` must be updated to use `firezone-gateway authenticate`, `firezone-gateway enable-service`, and `firezone-gateway disable-service` instead.

Authenticate a gateway and bring it into service using the consolidated `firezone-gateway` binary instead of the old separate `firezone` CLI.

```
$ firezone-gateway authenticate && firezone-gateway enable-service
```

Authenticate a gateway token directly from firezone-gateway without needing the separate firezone CLI.

```
$ firezone-gateway authenticate
```

Enable the gateway as a system service using the now-consolidated firezone-gateway CLI.

```
$ firezone-gateway enable-service
```

— Exact old and new commands given, migration path clear Gateway 1.6.2 (2026-09-24) changelog-20260924-986e97ff

THINNER COVERAGE BELOW

02 **iru-id SAN claim alias for device attestation** IMPROVED

35

Firezone now recognizes `iru-id` SAN claims in device attestation certificates as an alias for `kandji-id`.

— Names the claim but no usage context or config detail changelog-20260924-986e97ff

└──▷ **BREAKING ON UPGRADE**

! The `authenticate`, `enable-service`, and `disable-service` commands have moved from the `firezone` CLI binary into `firezone-gateway` — any scripts or service definitions calling `firezone authenticate`, `firezone enable-service`, or `firezone disable-service` must be updated to use `firezone-gateway authenticate`, `firezone-gateway enable-service`, and `firezone-gateway disable-service` instead.

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Firezone: Recognizes `iru-id` SAN claims in device attestation certificates as an alias for `kandji-id`.

└──▷ **TRY IT**

Authenticate a gateway token directly from firezone-gateway without needing the separate firezone CLI.

```
$ firezone-gateway authenticate
```

Enable the gateway as a system service using the now-consolidated firezone-gateway CLI.

```
$ firezone-gateway enable-service
```

Gateway 1.6.2 (2026-09-24) BLOG ↗ STABLE

Firezone Gateway 1.6.2 moves `authenticate`, `enable-service`, and `disable-service` commands into `firezone-gateway` itself.

└──▷ **TRY IT**

Authenticate a gateway and bring it into service using the consolidated `firezone-gateway` binary instead of the old separate `firezone` CLI.

```
$ firezone-gateway authenticate && firezone-gateway enable-service
```

└──▷ **BREAKING ON UPGRADE**

> ! The `authenticate`, `enable-service`, and `disable-service` commands have moved from the `firezone` CLI binary into `firezone-gateway` — any scripts or service definitions calling `firezone authenticate`, `firezone enable-service`, or `firezone disable-service` must be updated to use `firezone-gateway authenticate`, `firezone-gateway enable-service`, and `firezone-gateway disable-service` instead.

WAS THIS USEFUL? 👍 👎

Headscale

Feedback ⓘ v0.29.4 NOTES ↗ RANK

An open source, self-hosted implementation of the Tailscale control server

> Headscale v0.29.4 adds new Prometheus metrics for map-update and node-store activity, and reduces unnecessary peer-map rebuilds by keying the peer map by node ID.

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

01 **New Prometheus metrics for mapper and node-store** NEW

55

Adds `headscale_mapper_changes_dropped_total` and `headscale_ha_health_updates_total` metrics to track suppressed no-op map updates, plus `headscale_nodestore_snapshot_builds_total` to track peer-map snapshot rebuilds.

— Names three exact metric names but no usage guidance v0.29.4

02 **Peer map keyed by node ID to cut rebuild churn** IMPROVED

45

The peer map is now keyed by node ID and reused for writes that cannot change peer visibility, so routine map requests no longer trigger a full rebuild.

— Explains mechanism and effect but no config or metric to verify v0.29.4

v0.29.4 NOTES ↗ LATEST STABLE

Headscale v0.29.4 adds mapper and node-store metrics and reduces peer-map rebuild churn via node-ID keying.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.29.4
https://github.com/juanfont/headscale.git
# already have the repo? check out this version:
$ git checkout v0.29.4
```

Adds `headscale_mapper_changes_dropped_total` and `headscale_ha_health_updates_total` Prometheus metrics to track suppressed no-op map updates.

Adds `headscale_nodestore_snapshot_builds_total` Prometheus metric to track peer-map snapshot rebuilds.

> Peers the peer map by node ID and reuses it for writes that cannot change peer visibility, so routine map requests no longer trigger a full rebuild.

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Give feedback to the DNSFilter team so we can make more informed product decisions.
Powered by Canny.

> DNSFilter's macOS Roaming Client v2.4.6 adds IPv6 filtering, a DNS-over-HTTPS VPN fallback, and a new configuration.json architecture, alongside two breaking changes around macOS version support and certificate profiles.

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

01 Certificate profile consolidation and Root CA G1 requirement

85

BREAKING

macOS certificate profiles now install all DNSFilter roots — including DNSFilter Root CA G1 — in a single `Combined_Certificates.mobileconfig` push. v2.4.6 checks for this new root; devices without the updated profile show a Fix Installation Issues prompt in the menu bar until an admin pushes the updated `Combined_Certificates.mobileconfig` (or branded/white-label equivalent) via MDM, replacing the existing profile in place.

Push the updated combined certificate profile via MDM before auto-upgrade rolls out to clear the Fix Installation Issues prompt on managed devices.

> 📍 In your MDM console, navigate to the macOS Roaming Client install guide › Prerequisites section, download `Combined_Certificates.mobileconfig` (or your branded/white-label profile), and push it to all managed devices to replace the existing certificate profile in place.

— Names exact file and gives step-by-step MDM remediation.

09/23 September 23 - macOS Roaming Clie… changelog-20260924-4d627d69

02 DoH fallback during VPN sessions NEW

78

Adds the `ENABLE_DOH_FALLBACK=true` install-time flag to keep DNSFilter filtering active over DNS-over-HTTPS when a corporate VPN blocks the DNS-over-TLS port, so remote users stay filtered during VPN sessions.

Enable DNS-over-HTTPS fallback during VPN sessions so remote users stay filtered even when the VPN tunnel blocks DNS-over-TLS.

```env
ENABLE_DOH_FALLBACK=true
```

— Names exact install flag and the VPN scenario it solves.

09/23 September 23 - macOS Roaming Clie… changelog-20260924-4d627d69

03 **Travel Wi-Fi install-time configuration** IMPROVED 60

Travel Wi-Fi settings in `dns_agent.conf` now apply at CLI install time, and install-time configuration keys use the same Travel Wi-Fi names as the dashboard and documentation.

Pre-configure Travel Wi-Fi and site key at CLI install time using the now-documented key names in dns_agent.conf.

```yaml
ENABLE_DOH_FALLBACK=true
```

— Names config file and key alignment; CLI install path is actionable.

09/23 September 23 - macOS Roaming Clie… changelog-20260924-4d627d69

04 **IPv6 filtering support** NEW 60

Adds IPv6 filtering and protection for dual-stack and IPv6-only networks; administrators can set upstream resolution to IPv4-only, IPv6-only, or automatic.

— Explains the resolution modes but not the exact setting location.

09/23 September 23 - macOS Roaming Clie…

THINNER COVERAGE BELOW

05 **New configuration.json agent architecture** IMPROVED 55

Switches agent configuration from the legacy `daemon.conf` to a new `configuration.json` format; existing installations migrate automatically on upgrade.

— Names both file formats but migration is automatic, nothing to act on.

09/23 September 23 - macOS Roaming Clie…

06 **macOS 13 Ventura minimum requirement** BREAKING 50

The macOS Roaming Client now requires macOS 13 Ventura or later; devices running macOS 12 or earlier will not be able to run v2.4.6.

— States clear before/after version cutoff, no remediation steps given.

09/23 September 23 - macOS Roaming Clie… changelog-20260924-4d627d69

07 **Component version mismatch detection** NEW 45

The client now alerts when the menu bar app, system extension, and daemon are running different versions, with a user-facing fix prompt.

— Describes trigger and effect but no user action described.

09/23 September 23 - macOS Roaming Clie…

08 **Instant local domain and resolver updates** IMPROVED

25

Local domain and resolver changes made in the dashboard now apply immediately without a daemon restart.

— Bare before/after statement with no named surface.

09/23 September 23 - macOS Roaming Clie…

└──▷ **BREAKING ON UPGRADE**

! The macOS Roaming Client now requires macOS 13 Ventura or later; devices running macOS 12 or earlier will not be able to run v2.4.6.

! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices without the updated `Combined_Certificates.mobileconfig` profile will show a Fix Installation Issues prompt in the menu bar after upgrading, requiring an MDM push of the updated profile to clear.

! macOS 13 Ventura is now the minimum required OS version; devices running macOS 12 or earlier will not run the v2.4.6 client.

! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices that still have the previous certificate profile will show a Fix Installation Issues prompt until the updated `Combined_Certificates.mobileconfig` (or branded/white-label equivalent) is pushed via MDM.

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

└──▷ **USE IT**

Pre-configure Travel Wi-Fi and site key at CLI install time using the now-documented key names in dns_agent.conf.

dns_agent.conf

```
ENABLE_DOH_FALLBACK=true
```

└──▷ **BREAKING ON UPGRADE**

! The macOS Roaming Client now requires macOS 13 Ventura or later; devices running macOS 12 or earlier will not be able to run v2.4.6.

! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices without the updated `Combined_Certificates.mobileconfig` profile will show a Fix Installation Issues prompt in the menu bar after upgrading, requiring an MDM push of the updated profile to clear.

09/23 September 23 - macOS Roaming Client v2.4.6 Production (2026-09-23)

BLOG ↗ STABLE

macOS Roaming Client v2.4.6 adds IPv6 support, DoH VPN fallback, and a new configuration.json architecture

└──▷ **USE IT**

Enable DNS-over-HTTPS fallback during VPN sessions so remote users stay filtered even when the VPN tunnel blocks DNS-over-TLS.

```env
ENABLE_DOH_FALLBACK=true
```

Push the updated combined certificate profile via MDM before auto-upgrade rolls out to clear the `Fix Installation Issues` prompt on managed devices.

> 📍 In your MDM console, navigate to the macOS Roaming Client install guide › Prerequisites section, download `Combined_Certificates.mobileconfig` (or your branded/white-label profile), and push it to all managed devices to replace the existing certificate profile in place.

Adds `ENABLE_DOH_FALLBACK=true` install-time flag to keep DNSFilter filtering active over DNS-over-HTTPS when a corporate VPN blocks the DNS-over-TLS port.

Switches agent configuration from legacy `daemon.conf` to a new `configuration.json` format; existing installations migrate automatically on upgrade.

Travel Wi-Fi settings in `dns_agent.conf` now apply at CLI install time, and install-time configuration keys use the same Travel Wi-Fi names as the dashboard and documentation.

Adds IPv6 filtering and protection for dual-stack and IPv6-only networks; administrators can set upstream resolution to IPv4-only, IPv6-only, or automatic.

Adds component version mismatch detection: the client now alerts when the menu bar app, system extension, and daemon are running different versions, with a user-facing fix prompt.

`+2 more`

└──▷ **BREAKING ON UPGRADE**

- ! macOS 13 Ventura is now the minimum required OS version; devices running macOS 12 or earlier will not run the v2.4.6 client.
- ! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices that still have the previous certificate profile will show a `Fix Installation Issues` prompt until the updated `Combined_Certificates.mobileconfig` (or branded/white-label equivalent) is pushed via MDM.

WAS THIS USEFUL? 👍 👎

Cato Networks

Feedback ⓘ

Securing what your Bedrock AgentCore agents actually do (2026-09-24) BLOG ↗

RANK

Cato SASE Cloud converges SD-WAN, security, global backbone, and remote access into a global cloud-native service. WE ARE SASE™

> Cato launched AI Security for Agents, adding read-only discovery and inline inspection of AI agents running on Amazon Bedrock AgentCore, alongside broadened coverage across major agent-hosting platforms.

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

01 **Inline inspection of agent-model exchanges** NEW 63

Adds inline inspection of agent-model exchanges at four observable stages — prompt, tool call, tool response, and output — to detect prompt injection and other unwanted activity before the model acts.

— Names the four inspection stages but no activation path given.

Securing what your Bedrock AgentCore ag…

THINNER COVERAGE BELOW

02 **Read-only discovery of Bedrock AgentCore agents** NEW 58

Cato AI Security for Agents adds read-only discovery of every agent running on Amazon Bedrock AgentCore, surfacing each agent's instructions, model, and connected tools without installing anything.

— Describes mechanism and scope but no setup steps or config given.

Securing what your Bedrock AgentCore ag…

03 **AI Security for Agents platform coverage expanded** NEW 51

Extends Cato AI Security for Agents coverage to Amazon Bedrock (Agents Classic and AgentCore), Microsoft Foundry, Microsoft Copilot Studio, ChatGPT (GPTs and Agents), Amazon Quick, and developer-hosted agent endpoints.

— Lists supported platforms but no mechanism or how-to detail.

Securing what your Bedrock AgentCore ag…

Securing what your Bedrock AgentCore agents actually do (2026-09-24) BLOG ↗

LATEST STABLE

> Cato AI Security for Agents adds read-only discovery and inline inspection of Amazon Bedrock AgentCore agents across prompt, tool call, tool response, and output stages.
>
> Adds read-only discovery of every agent running on Amazon Bedrock AgentCore, surfacing each agent's instructions, model, and connected tools without installing anything.
>
> Adds inline inspection of agent-model exchanges at four observable stages — prompt, tool call, tool response, and output — to detect prompt injection and other unwanted activity before the model acts.
>
> Extends Cato AI Security for Agents coverage to Amazon Bedrock (Agents Classic and AgentCore), Microsoft Foundry, Microsoft Copilot Studio, ChatGPT (GPTs and Agents), Amazon Quick, and developer-hosted agent endpoints.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Sep 22nd, 2026 NOTES ↗ RANK

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

> Compyl added a new Box integration for compliance evidence collection and significantly deepened its GitHub and Okta integrations, while also making custom field values searchable across core workspace grids.

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

01 **GitHub integration expanded with security signals** IMPROVED 75

The GitHub integration now covers Dependabot alerts, secret scanning alerts, code scanning alerts, organization security settings (including 2FA status), outside collaborators, and rulesets.

— Lists every new data surface pulled in, no config detail Sep 22nd, 2026

02 **Okta integration expanded with identity signals** IMPROVED 75

The Okta integration now covers application entitlements, MFA enrollment status, group rules, admin role assignments, and user lifecycle events.

— Names each new data point, but no mechanism detail Sep 22nd, 2026

03 **Box integration for compliance evidence** NEW 70

Compyl adds Box as a new integration, collecting users, groups, files, folders, collaborations, shared links, legal holds, retention policies, classifications, and enterprise events for continuous compliance evidence.

— Names data types collected but no setup steps or limits Sep 22nd, 2026

04 **Searchable custom fields in grid views** IMPROVED 60

Custom field values are now searchable and filterable directly in grid views across Risks, Vendors, Incidents, Contracts, Assets, and Audits.

— Names affected views clearly, no exact UI steps given Sep 22nd, 2026

THINNER COVERAGE BELOW

05 **MongoDB Atlas integration expanded** IMPROVED 10

Compyl expanded its MongoDB Atlas integration coverage, though the specific new data points collected are not detailed.

— Mentioned only in summary with no concrete detail Sep 22nd, 2026

└──▷ ALSO FROM THESE RELEASES

Sep 22nd, 2026 NOTES ↗ LATEST STABLE

Compyl adds Box integration and expands GitHub, Okta, and MongoDB Atlas coverage; custom fields now searchable in grid views.

Adds Box as a new integration, collecting users, groups, files, folders, collaborations, shared links, legal holds, retention policies, classifications, and enterprise events for continuous compliance evidence.

Expands the GitHub integration to cover Dependabot alerts, secret scanning alerts, code scanning alerts, organization security settings (including 2FA status), outside collaborators, and rulesets.

Expands the Okta integration to cover application entitlements, MFA enrollment status, group rules, admin role assignments, and user lifecycle events.

Makes custom field values searchable and filterable directly in grid views across Risks, Vendors, Incidents, Contracts, Assets, and Audits.

WAS THIS USEFUL? 👍 👎

◆ Third-Party & Vendor Risk (TPRM)

Conveyor

Feedback ⓘ

🌐 Work on Non-English Word Questionnaires in English NOTES ↗ RANK

Conveyor delivers AI for high-stakes customer requests like security reviews, vendor assessments, and requests for proposal (RFP).

> Conveyor added support for working on non-English Word questionnaires entirely in English before exporting back to the original language.

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

01 Translation workflow for non-English Word questionnaires NEW

60

Non-English Word questionnaires can now be imported through the Questionnaires module, translated so reviewers can work entirely in English, and then exported in either the original language or English.

When you receive a security questionnaire in a foreign language, import it through the Questionnaires module to work on it entirely in English, then export back in the original language for the requester.

> 📍 1. Go to the Questionnaires module. 2. Upload your non-English Word questionnaire as you normally would. 3. Review and answer all questions in the English-translated view. 4. When finished, choose to export in the original language or keep it in English.

— Names module and workflow steps but no limits or formats.

🌐 Work on Non-English Word Questionnai…

└──▷ **ALSO FROM THESE RELEASES**

. . . .

🌐 Work on Non-English Word Questionnaires in English NOTES ↗ LATEST

STABLE

Conveyor now translates non-English Word questionnaires to English for review, then exports in the original language or English.

└──▷ **HOW TO FIND IT**

When you receive a security questionnaire in a foreign language, import it through the Questionnaires module to work on it entirely in English, then export back in the original language for the requester.

📍 1. Go to the Questionnaires module. 2. Upload your non-English Word questionnaire as you normally would. 3. Review and answer all questions in the English-translated view. 4. When finished, choose to export in the original language or keep it in English.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗ RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

> Runlayer's biggest shipments this cycle are DPoP-bound agent tokens enforced at the MCP proxy and a new Salesforce connector, alongside calendar-based LLM budget windows, expanded trigger management (a Gmail label trigger plus MCP-based trigger CRUD), and a long tail of smaller changes across Slack agent behavior, audit logging, skills, and AI Watch scan exit codes.

└──▷ **WHAT SHIPPED · 21 FEATURES** most completely described first ? what's the number?

01 **DPoP-bound agent tokens** NEW 80

Agent-account tokens can be bound to DPoP keys at issuance and enforced at the MCP proxy, with a per-agent `require DPoP` flag and admin UI; once enabled, the proxy rejects any token presented without a valid DPoP proof, preventing replay of intercepted tokens.

Enforce DPoP binding on a high-privilege agent so its tokens cannot be replayed if intercepted.

> 📍 In the Runlayer console, open the agent › Settings › Security, enable Require DPoP, then re-issue the agent-account token with a DPoP key bound at issuance. The MCP proxy will reject any token presented without a valid DPoP proof.

— Mechanism, enforcement point and UI steps to enable are all given.

changelog-20260924-c2abe11d

02 **Salesforce connector** NEW 75

Adds Salesforce connector to the catalog, carrying the Salesforce server's own tool schemas so agents and triggers call its tools with the correct arguments; the setup guide documents the required OAuth scope and `INVALID_SCOPES` recovery.

— Names OAuth scope requirement and a specific recoverable error code.

changelog-20260924-c2abe11d

03 **OpenAI-Compatible provider routing for Gemini** IMPROVED 70

Runlayer can route a Gemini model through an OpenAI-Compatible provider when an agent already targets an OpenAI-style endpoint: in the console under Providers > Add Provider > OpenAI-Compatible, set the base URL to

`https://generativelanguage.googleapis.com` and paste a Google API key, which Runlayer sends as a Google API key rather than a Bearer token.

Route a Gemini model through an OpenAI-Compatible provider when you want to use Google Gemini from an agent that already targets an OpenAI-style endpoint.

> 📍 In the Runlayer console, go to Providers › Add Provider › OpenAI-Compatible, set the base URL to `https://generativelanguage.googleapis.com`, and paste your Google API key. Runlayer will send it as a Google API key rather than a Bearer token. Then select the provider when creating or editing your agent.

— Exact console path, base URL, and key-handling behavior given.

changelog-20260924-c2abe11d

---

04 **AI Watch scan exit status change** BREAKING 68

AI Watch scan exit status now exits 0 with a Submitted; server marked surfaces incomplete warning (naming each surface and reason) for scans the server accepted but could not fully attribute, instead of exiting 3 as Could not submit.

— Exact before/after exit codes and message text given. changelog-20260924-c2abe11d

---

05 **google-docs edit_table set_cells support** NEW 60

Adds `google-docs` `edit_table` tool support for setting multiple table cells in one call via `set_cells`.

— Names exact tool and parameter for direct use. changelog-20260924-c2abe11d

THINNER COVERAGE BELOW

---

06 **Audit log expansion** IMPROVED 57

Adds OAuth token redaction in audit log details; adds Agent Discovered as a named audit event type for agents found by AI Watch device scans, including back-filling previously recorded Unknown Event rows; audits managed Slack provisioning; and agent-run tool-call audit events now carry the agent and run IDs.

— Names four distinct audit additions with backfill detail, no query surface given.

changelog-20260924-c2abe11d

---

07 **max_completion_tokens for Azure Foundry gpt-5 agents** BREAKING 55

Adds `max_completion_tokens` in place of the legacy `max_tokens` parameter for Azure Foundry agents on gpt-5-family deployments.

— Names exact parameter and deployment family, no migration guidance given.

changelog-20260924-c2abe11d

08 **Gmail label trigger and MCP trigger management** NEW

55

Adds a Gmail label-added poll trigger that fires an agent when a label is added to a Gmail message, and adds trigger management over MCP so agents can create, list, and delete agent triggers through Runlayer MCP tools.

– Names both trigger surfaces but not the exact MCP tool names.

changelog-20260924-c2abe11d

09 **Slack agent interaction improvements** IMPROVED

50

Agents can continue following a Slack thread after being mentioned without requiring another mention; long-running agent steps now show progress on a fresh Slack card without a warning icon when a step takes more than a few minutes; and the agent channel intro message is now off by default, configurable per agent.

– Three distinct Slack behavior changes described without exact config keys.

changelog-20260924-c2abe11d

10 **Per-tool enrollment diff in security-setting enrollment** IMPROVED

49

Adds a per-tool enrollment diff in the recommended security-setting enrollment flow, showing the effective per-tool diff with scan-direction safety checks before applying.

– Explains what the diff shows, not the exact UI path. changelog-20260924-c2abe11d

11 **Skill zip import and rm -rf detection** NEW

48

Adds zip skill import to create a multi-file skill by importing a single zip file, and extends skill scanning to catch combined short-flag recursive deletes such as `rm -rf`.

– Names the import mechanism and the specific flag pattern detected.

changelog-20260924-c2abe11d

12 **Run schedules on demand** NEW

45

Adds Run schedules on demand : run a schedule immediately and retry failed scheduled runs from the schedule list.

– Names UI location and both actions available there. changelog-20260924-c2abe11d

13 **Run continuation after approval** IMPROVED

43

Adds Runs continue after approval for MCP- and REST-triggered agent runs, and allows continuing your own run from any surface.

– Names the trigger types affected but not an exact endpoint. changelog-20260924-c2abe11d

14 **Playground Stop button** NEW

42

Adds a Playground Stop button that cancels the active run.

— Names the exact control and its action. changelog-20260924-c2abe11d

15 **Inline video artifact playback** IMPROVED 40

Adds inline video artifact playback in the run view, with the artifact size limit raised to 25 MB.

— Exact size limit given, little further mechanism. changelog-20260924-c2abe11d

16 **View as permission group for Super Admins** NEW 39

Adds a View as permission group capability allowing Super Admins to view the workspace as any permission group.

— Clear capability named, no exact navigation path given. changelog-20260924-c2abe11d

17 **Sessions timeline with skipped Slack messages** IMPROVED 38

The agent Sessions list now shows messages the agent chose to skip alongside runs, and search covers both skipped messages and runs.

— Describes behavior but no exact UI path or field names. changelog-20260924-c2abe11d

18 **Calendar-based budget windows for LLM usage rules** NEW 35

Adds calendar-based budget windows to LLM gateway usage rules, enabling monthly or other calendar-period caps.

— States the capability but no config key or exact cap values.
changelog-20260924-c2abe11d

19 **Login-host badge on managed devices** NEW 35

Adds a login-host badge on managed devices shown when a user's login host differs from the MDM-managed host.

— Names the trigger condition but no further mechanism. changelog-20260924-c2abe11d

20 **Model picker pricing and context data** IMPROVED 35

Model picker hover cards now show pricing and context data from the LiteLLM catalog.

— Names the data source but little mechanism beyond display. changelog-20260924-c2abe11d

21 **Vercel integration via OAuth broker** NEW 26

Adds Vercel integration available through the OAuth broker.

— Bare mention with no setup or scope detail. changelog-20260924-c2abe11d

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Runlayer adds Claude Opus 5.5, GPT-6 models, DPoP-bound agent tokens, calendar budget windows, MCP trigger management, and a Gmail label trigger across this release cycle.

└──▷ HOW TO FIND IT

Route a Gemini model through an OpenAI-Compatible provider when you want to use Google Gemini from an agent that already targets an OpenAI-style endpoint.

> 📍 In the Runlayer console, go to Providers › Add Provider › OpenAI-Compatible, set the base URL to `https://generativelanguage.googleapis.com`, and paste your Google API key. Runlayer will send it as a Google API key rather than a Bearer token. Then select the provider when creating or editing your agent.

Enforce DPoP binding on a high-privilege agent so its tokens cannot be replayed if intercepted.

> 📍 In the Runlayer console, open the agent › Settings › Security, enable `Require DPoP`, then re-issue the agent-account token with a DPoP key bound at issuance. The MCP proxy will reject any token presented without a valid DPoP proof.

Adds DPoP-bound agent tokens: agent-account tokens can be bound to DPoP keys at issuance and enforced at the MCP proxy, with a per-agent `require DPoP` flag and admin UI.

Adds calendar-based budget windows to LLM gateway usage rules, enabling monthly or other calendar-period caps.

Adds `max_completion_tokens` in place of the legacy `max_tokens` parameter for Azure Foundry agents on gpt-5-family deployments.

Adds a Gmail label-added poll trigger that fires an agent when a label is added to a Gmail message.

Adds Trigger management over MCP: create, list, and delete agent triggers through Runlayer MCP tools.

`+21 more`

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

AI-Infra-Guard

Feedback ⓘ v4.6.3 NOTES ↗ RANK

A full-stack AI Red Teaming platform securing AI ecosystems via Agent Scan, Skills Scan, MCP scan, AI Infra scan and LLM jailbreak evaluation.

> AI-Infra-Guard's v4.6.3 release focuses on reliability improvements to DeepTeam test execution and API Checker model fingerprinting.

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

01 **Mid-run disk persistence for DeepTeam** IMPROVED 35

DeepTeam now writes finished test cases to disk during a run rather than holding them only in memory, reducing the risk of losing completed results if a run is interrupted.

— Explains mechanism but no config or flag named v4.6.3

02 **API Checker fingerprint sample refilling** IMPROVED 30

API Checker improves model identification by refilling failed fingerprint samples, making model detection more reliable.

— States what changed but lacks specifics on scope or metrics v4.6.3

v4.6.3 NOTES ↗ LATEST STABLE

DeepTeam now flushes finished test cases to disk mid-run; API Checker gets more reliable model fingerprinting via sample refilling.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v4.6.3 https://github.com/Tencent/AI-Infra-Guard.git
# already have the repo? check out this version:
$ git checkout v4.6.3
```

Adds mid-run disk persistence for DeepTeam: finished test cases are written to disk during a run rather than held only in memory.

Improves API Checker model identification by refilling failed fingerprint samples for more reliable detection.

WAS THIS USEFUL? 👍 👎

NVIDIA SkillSpector

Feedback ⓘ v2.12.0 NOTES ↗ RANK

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, security risks, prompt injection, data exfiltration, and supply-chain risks in Claude Code, Codex, and MCP skills before you install them.

> SkillSpector v2.12.0 adds a CI-friendly `--fail-on-findings` exit gate, SC10 detection of noncanonical dependency sources across major package managers, native OpenCode integration via `SKILLSPECTOR_PROVIDER=opencode_cli` and a `.opencode/` extension, and support for scanning GitHub `/tree/` subdirectories, alongside a run of smaller reporting, provider, and detection refinements.

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

01 **Fail-on-findings CI gate** NEW 90

Adds `skillspector scan --fail-on-findings` flag that exits with code 1 when any active finding is present, including findings below the default risk-score threshold, across single-skill, recursive, and MCP registry scans.

Fail CI immediately when a skill scan surfaces any active finding, useful as a hard gate in a pull-request pipeline.

```
$ skillspector scan ./my-skill/ --fail-on-findings
```

— Exact flag and exit-code behaviour with a runnable command v2.12.0

02 **OpenCode provider and native extension** NEW 90

Adds `SKILLSPECTOR_PROVIDER=opencode_cli` to run semantic analysis through a local OpenCode login, requiring exactly OpenCode 1.18.31 and failing closed on auth, version, policy, or empty-output failures. Also adds a native `/skillspector` command and `skillspector_scan` tool to OpenCode via the repository-provided `.opencode/` extension, installed by copying `.opencode/` from a checkout and resolved via `SKILLSPECTOR_BIN`, a worktree `.venv`, or `PATH`.

— Full env var, version pin, command, tool and install paths named v2.12.0

03 **Static analysis time budget env var** NEW 80

Adds `SKILLSPECTOR_MAX_STATIC_ANALYSIS_SECONDS_PER_ARTIFACT` environment variable to configure static-pattern and YARA time allowance per artifact; default raised from 30 to 300 seconds.

— Named env var with old/new defaults, no example v2.12.0

04 **SC10 dependency-source detection** NEW 75

Adds SC10 detection of noncanonical or unresolved dependency sources across npm, Yarn, pip, Poetry, Maven, and Cargo configuration, shell scripts, generated heredocs, and actionable shell fences in `SKILL.md` / `README.md`.

— Lists every ecosystem covered but no example run v2.12.0

05 **GitHub tree URL scanning** NEW 75

Supports HTTPS GitHub `/tree/` URLs to select a repository subdirectory as scan input, resolving the longest advertised branch/tag name including refs containing slashes.

Scan only a specific subdirectory of a GitHub repository, such as a single skill inside a monorepo.

```
$ skillspector scan https://github.com/org/skills-monorepo/tree/main/skills/my-skill
```

— Mechanism explained plus a runnable command v2.12.0

06 **LLM provenance metadata in JSON reports** NEW 75

Adds `metadata.llm_provenance` (schema version 1) to JSON reports, recording sanitized model/package/source identity, configured vs. effective providers, and sampling/reasoning controls; `SKILLSPECTOR_BUILD_REVISION` can supply source revision metadata.

— Names field, schema version and env var, no worked example v2.12.0

07 **Compact LLM prompt labels** NEW 65

Adds `SKILLSPECTOR_COMPACT_PROMPTS=1` (also `true` / `yes`) to enable shorter unpadded LLM prompt line labels such as `L1:` instead of `L01:`.

— Named env var and accepted values, minor cosmetic scope v2.12.0

08 **Model registry configuration** IMPROVED 65

Registers `gemini-3.5-flash` model budgets for the root and OpenAI-compatible registries, enabling use through the existing `openai` provider; CLI-backed providers now also honor `SKILLSPECTOR_MODEL_REGISTRY` for per-model context and output-token limits.

— Names model, provider and env var but no example usage v2.12.0

09 **Missing vs unresolved reference distinction** IMPROVED 60

Distinguishes `reference_missing` from `reference_unresolved` in MCP installation decisions, allowing scans whose only caveats are missing references to

proceed when all discovered files were inspected.

— Named states and behaviour change, no direct action for user v2.12.0

THINNER COVERAGE BELOW

10 **JSON and SARIF column reporting change** IMPROVED 55

JSON findings retain each occurrence's own zero-based columns; SARIF emits one-based Unicode-codepoint columns for active and suppressed findings.

— Concrete format detail but consumer-facing only v2.12.0

11 **Terminal scan progress streaming** IMPROVED 50

Non-verbose interactive terminal scans now stream graph execution progress, completed analyzer rules, and a discovered-file tree to standard error without affecting machine-readable standard output.

— Explains behaviour but no flag or command to invoke v2.12.0

12 **GHSA/OSV MODERATE severity remapped to MEDIUM** BREAKING 50

GHSA/OSV `MODERATE` severity now maps to `MEDIUM` instead of the previous `LOW`, potentially increasing affected supply-chain risk scores.

— Clear before/after mapping but no migration guidance given v2.12.0

13 **TP4 semantic analysis on fenced code** IMPROVED 40

Extends TP4 semantic analysis to non-empty, closed Markdown/text fences with recognized executable-language labels.

— Brief mechanism note, no scope numbers or example v2.12.0

14 **Reflective Python module and bytecode detection** NEW 30

Detects reflective Python module lookups and shipped bytecode inside Python environments.

— Bare description with no mechanism or example v2.12.0

v2.12.0 NOTES ↗ LATEST STABLE

SkillSpector v2.12.0 adds `--fail-on-findings`, SC10 supply-chain detection, opencode_cli provider, TP4 Markdown fences, and GitHub tree scanning.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.12.0
https://github.com/NVIDIA/SkillSpector.git
# already have the repo? check out this version:
$ git checkout v2.12.0
```

└──▷ **TRY IT**

Fail CI immediately when a skill scan surfaces any active finding, useful as a hard gate in a pull-request pipeline.

```
$ skillspector scan ./my-skill/ --fail-on-findings
```

Scan only a specific subdirectory of a GitHub repository, such as a single skill inside a monorepo.

```
$ skillspector scan https://github.com/org/skills-
monorepo/tree/main/skills/my-skill
```

Adds `skillspector scan --fail-on-findings` flag that exits with code 1 when any active finding is present, including findings below the default risk-score threshold, across single-skill, recursive, and MCP registry scans.

Adds `SKILLSPECTOR_MAX_STATIC_ANALYSIS_SECONDS_PER_ARTIFACT` environment variable to configure static-pattern and YARA time allowance per artifact; default raised from 30 to 300 seconds.

Adds `SKILLSPECTOR_COMPACT_PROMPTS=1` (also `true` / `yes`) to enable shorter unpadded LLM prompt line labels such as `L1:` instead of `L01:`.

Adds `SKILLSPECTOR_PROVIDER=opencode_cli` to run semantic analysis through a local OpenCode login, requiring exactly OpenCode 1.18.31 and failing closed on auth, version, policy, or empty-output failures.

Adds a native `/skillspector` command and `skillspector_scan` tool to OpenCode via the repository-provided `.opencode/` extension; installed by copying `.opencode/` from a checkout, resolved via `SKILLSPECTOR_BIN`, a worktree `.venv`, or `PATH`.

+11 more

WAS THIS USEFUL? 👍 👎

ai-safe2-framework

Feedback ⓘ

2026-09-23_NEXUS_v0.4-Safer_Payment_Authority_for_Agents NOTES ↗ RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> NEXUS v0.4 introduces eight named reference profiles that give AI agents deterministic, fail-closed authority decisions across major payment protocols, backed by an adversarial test corpus.

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

01 **Eight payment authority adapter profiles for agents** NEW

75

NEXUS v0.4 adds eight reference profiles covering payment authority evidence for AI agents: `SafePay Exact` for structural x402 v2 exact EVM/USDC evaluation, `SafePay Verified` for authenticated x402 facilitator or chain evidence with an explicit verifier boundary, `SafePay Variable` for maximum-versus-actual charge enforcement and reservation release, `SafePay Escrow` for deposit, release, refund, and mediated-finality state handling, Mandate Bridge for AP2 v0.2 mandate and receipt evidence, Card Trust Bridge for Visa TAP request recognition without overstating payment authorization, Agent Token Bridge for authenticated Mastercard AP4M verifier evidence, and Portable Delegation Bridge for KYA-OS identity and delegated-authority evidence.

— Names all eight profiles and their protocol scope, but no config or invocation shown

2026-09-23_NEXUS_v0.4-Safer_Payment_Aut…

THINNER COVERAGE BELOW

02 **Deterministic authority outcomes with fail-closed guardrails**

NEW

55

Introduces three deterministic assessment outcomes — Supported, Contradicted, and Unverifiable — giving agents stable machine-readable authority decisions instead of permission inferred from prose. Provides explicit allowlists, assurance ceilings, and fail-closed results across all eight adapter profiles.

— Explains outcome model and guardrails but no concrete config surface

2026-09-23_NEXUS_v0.4-Safer_Payment_Aut…

03 **Adversarial test corpus for payment adapters** NEW

55

Includes an adversarial test corpus covering substitution attacks, malformed evidence, truthy-value bypasses, replay behavior, scope mapping, privacy checks, and RailGuard conformance.

— Names specific attack categories tested but not how to run the corpus

2026-09-23_NEXUS_v0.4-Safer_Payment_Aut…

2026-09-23_NEXUS_v0.4-Safer_Payment_Authority_for_Agents NOTES ↗ LATEST

STABLE

NEXUS v0.4 adds eight named payment authority adapter profiles for AI agents, covering x402, AP2, Visa TAP, Mastercard AP4M, and KYA-OS delegation evidence.

└──▷ GET THIS VERSION

```
$ git clone --branch 2026-09-23_NEXUS_v0.4-
Safer_Payment_Authority_for_Agents
https://github.com/CyberStrategyInstitute/ai-safe2-framework.git
# already have the repo? check out this version:
$ git checkout 2026-09-23_NEXUS_v0.4-
Safer_Payment_Authority_for_Agents
```

Adds `SafePay Exact` reference profile for structural x402 v2 exact EVM/USDC evaluation.

Adds `SafePay Verified` reference profile for authenticated x402 facilitator or chain evidence with an explicit verifier boundary.

Adds `SafePay Variable` reference profile for maximum-versus-actual charge enforcement and reservation release.

Adds `SafePay Escrow` reference profile for deposit, release, refund, and mediated-finality state handling.

Adds `Mandate Bridge` reference profile for AP2 v0.2 mandate and receipt evidence.

+6 more

WAS THIS USEFUL? 👍 👎

Manifold Security

Feedback ⓘ announced · Sep 23, 2026

BLOG ↗ RANK

Manifold Security is an AI detection and response platform that monitors what autonomous AI agents do at runtime on enterprise endpoints and blocks risky actions.

> Manifold Security's launch introduces an `Ask` policy response that pauses AI coding agent actions for developer consent, delivered via inline prompts or native OS dialogs across eight supported agents, alongside an `Alert` mode for observing new rules before enforcing them.

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

01 **Ask policy for pausing agent actions on consent** NEW

80

Introduces an `Ask` policy response that halts a coding agent's action mid-task and waits for the developer to approve or reject it before proceeding, treating silence as a rejection. Enforcement of `Ask`, `Alert`, `Allow`, and `Deny` extends to all supported agents — Claude Code, Cursor, Gemini CLI, GitHub Copilot, and Antigravity, which surface the prompt inline via each agent's own permission system, plus Codex, Devin, and Kimi, which receive a native OS consent dialog on the developer's machine.

— Names the policy, mechanism, and every supported agent and delivery path

`launch-20260923-b4e89370`

02 **Alert policy for observation-mode rule tuning** NEW

61

The `Alert` policy response lets a new rule run in observation mode: the action proceeds, the developer is notified in-session or on the desktop, and the event is recorded centrally, enabling hit-rate tuning before committing the rule to `Ask` or `Deny`.

— Explains mechanism and purpose but no config surface named `launch-20260923-b4e89370`

> announced · Sep 23, 2026 BLOG ↗
>
> Manifold Security adds an `Ask` policy response that pauses agent actions for developer consent across all supported coding agents.
>
> Adds an `Ask` policy response that halts a coding agent's action mid-task and waits for the developer to approve or reject it before proceeding — silence is treated as a rejection.
>
> Extends `Ask`, `Alert`, `Allow`, and `Deny` policy enforcement to all supported agents: Claude Code, Cursor, Gemini CLI, GitHub Copilot, and Antigravity (via each agent's

own permission prompt), plus Codex, Devin, and Kimi (via a native OS consent dialog).

Introduces a two-path delivery model: agents with their own approval system surface the Manifold prompt inline in the session; agents without one receive a native OS dialog on the developer's machine.

The `Alert` policy response lets a new rule run in observation mode — action proceeds, the developer is notified in-session or on the desktop, and the event is recorded centrally — enabling hit-rate tuning before committing to `Ask` or `Deny`.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

The free and flexible app for your private thoughts.

└──▷ HOW TO FIND IT

Open a Markdown file that lives outside your vault — useful for reviewing reports or notes stored in a different directory without moving them.

📍 In Obsidian, open the Command Palette, search for 'Open file from outside the vault...', select it, then browse to any file on your computer.

Set Obsidian as the default Markdown viewer so double-clicking any .md file on your OS opens it directly in Obsidian.

📍 Download the latest Obsidian installer, run it, then right-click any .md file in your OS file manager, choose `Open with › Obsidian`, and select `Always use this app`.

└──▷ BREAKING ON UPGRADE

! Base font size is no longer stored in `appearance.json` and will no longer sync via Obsidian Sync; each device will need its font size set independently.

WAS THIS USEFUL? 👍 👎

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Sep 24, 2026 CHANGELOG ↗

RANK

Connect, protect, and build everywhere.

> Cloudflare's biggest addition this window is per-pull-request Worker Previews with isolated URLs and observability, alongside a new Gateway traffic-destination selector, expanded browser-binding session controls, guided Cloudflare Mesh onboarding, and account-level Images transformation analytics.

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

01 Traffic Destination selector in Gateway policies NEW

75

Adds the `net.offramp.type` API field, exposed as the Traffic Destination selector in Gateway HTTP and Network policies, with values `internet`, `cloudflare_wan`, `cloudflare_tunnel`, `device_client`, and `mesh` to target specific off-ramp methods so different policies can apply to traffic leaving via different paths.

Write a Gateway Network policy that applies stricter controls only to traffic leaving via Cloudflare Tunnel, while letting public Internet traffic follow a separate rule.

```
$ # In the Cloudflare dashboard or via the Zero Trust API, add a
Network policy condition:
# Selector: Traffic Destination | Operator: is | Value:
cloudflare_tunnel
# API field equivalent:
net.offramp.type == "cloudflare_tunnel"
```

Apply a stricter Gateway HTTP policy only to traffic headed for the public Internet, leaving tunnel-routed traffic under a different rule set.

```text
# In a Gateway HTTP policy expression:
net.offramp.type == "internet"
```

— Names field, selector, and all five values with a policy example

changelog-20260924-cfc65a33

THINNER COVERAGE BELOW

02 **Guided onboarding for Cloudflare Mesh** NEW 53

Adds guided onboarding for Cloudflare Mesh participants directly from the dashboard, with deployment options for Linux, Kubernetes, Docker Compose, and Docker CLI, plus a unified node-and-device management table.

— Names deployment targets but only a UI starting point changelog-20260924-cfc65a33

03 **Transformation analytics in Images** NEW 44

Introduces account-level transformation analytics in Images, showing request traffic split by URL-based vs. Images binding transformations, top zones, configurations, origin hosts, and Worker scripts.

— Describes dashboard breakdowns but no navigation or command given

changelog-20260924-98245c68

└──▷ ALSO FROM THESE RELEASES

Add a 32-argument concat() header-rewrite rule to bundle request context into a single origin header, using the Rulesets API on the `http_request_late_transform` phase.

```
$ curl --request POST \

"https://api.cloudflare.com/client/v4/zones/$ZONE_ID/rulesets/$RULESET
_ID/rules" \
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --data '{
    "ref": "add_request_context_header",
    "description": "Add request context for the origin",
    "expression": "true",
    "action": "rewrite",
    "action_parameters": {
      "headers": {
        "X-Request-Context": {
          "operation": "set",
          "expression": "concat(\"ip=\", to_string(ip.src),
\";country=\", ip.src.country, \";host=\", http.host, \";method=\",
http.request.method, \";path=\", http.request.uri.path, \";query=\",
http.request.uri.query, \";ray-id=\", cf.ray_id, \";asn=\",
to_string(ip.src.asnum), \";user-agent=\", http.user_agent)"
        }
      }
    }
  }'
```

changelog · Sep 24, 2026 CHANGELOG ↗ LATEST STABLE

Cloudflare: Introduces Worker Previews, letting each pull request run in an isolated environment with its own code, URL, bindings, logs, metrics, and traces, configured via the `previews` block in the Wrangler configuration file and deployed with `npx wrangler preview` (requires Wrangler 4.135.0 or later).

.
.
.

└──▷ **TRY IT**

Write a Gateway Network policy that applies stricter controls only to traffic leaving via Cloudflare Tunnel, while letting public Internet traffic follow a separate rule.

```
$ # In the Cloudflare dashboard or via the Zero Trust API, add a
Network policy condition:
# Selector: Traffic Destination | Operator: is | Value:
cloudflare_tunnel
# API field equivalent:
net.offramp.type == "cloudflare_tunnel"
```

Deploy a Worker Preview for a feature branch so reviewers get an isolated, production-like URL without affecting production or other previews.

```
$ npx wrangler preview
```

Add a 32-argument concat() header-rewrite rule to bundle request context into a single origin header, using the Rulesets API on the `http_request_late_transform` phase.

```
$ curl --request POST \

"https://api.cloudflare.com/client/v4/zones/$ZONE_ID/rulesets/$RULES
ET_ID/rules" \
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --data '{
    "ref": "add_request_context_header",
    "description": "Add request context for the origin",
    "expression": "true",
    "action": "rewrite",
    "action_parameters": {
      "headers": {
        "X-Request-Context": {
          "operation": "set",
          "expression": "concat(\"ip=\", to_string(ip.src),
\";country=\", ip.src.country, \";host=\", http.host, \";method=\",
http.request.method, \";path=\", http.request.uri.path, \";query=\",
http.request.uri.query, \";ray-id=\", cf.ray_id, \";asn=\",
to_string(ip.src.asnum), \";user-agent=\", http.user_agent)"
        }
      }
    }
  }'
```

Introduces account-level transformation analytics in Images, showing request traffic split by URL-based vs. Images binding transformations, top zones, configurations, origin hosts, and Worker scripts.

changelog · Sep 24, 2026 CHANGELOG ↗ STABLE

Cloudflare adds Traffic Destination selector, Worker Previews, private MCP support, browser binding session APIs, and more

.
.
.

└──▷ **USE IT**

Apply a stricter Gateway HTTP policy only to traffic headed for the public Internet, leaving tunnel-routed traffic under a different rule set.

```text
# In a Gateway HTTP policy expression:
net.offramp.type == "internet"
```

Deploy a Worker Preview for a feature branch so reviewers get an isolated, production-like URL without touching production.

```
$ npx wrangler preview
```

Route browser-binding requests for an internal hostname through an outbound Worker that adds auth, instead of exposing the service publicly.

```javascript
const connection = await BROWSER.acquire({
  outboundByHost: {
    "private.example.test": env.OUTBOUND
  }
});
const session = await BROWSER.connectSession(sessionId);
```

Adds `net.offramp.type` API field as the Traffic Destination selector in Gateway HTTP and Network policies, with values `internet`, `cloudflare_wan`, `cloudflare_tunnel`, `device_client`, and `mesh` to target specific off-ramp methods.

Adds acquire() and launch() methods to browser bindings with an `outboundByHost` option to route per-hostname requests through another Worker, plus connectSession(sessionId) and a nested `devtools` binding for typed DevTools and session-management operations.

Adds guided onboarding for Cloudflare Mesh participants from the dashboard, with deployment options for Linux, Kubernetes, Docker Compose, and Docker CLI, plus a unified node-and-device management table.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 24, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue