

# Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 16, 2026 · EVERY DAY

EDITIONS **tail** | grep | head | ls | diff | uniq

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 112 changes](#)Do you prefer **scroll**?[▶ // CONTENTS](#) [SHOW](#)

36 tools

[▼ // THE BRIEF](#) [HIDE](#)

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

Three unrelated things stand out. Zenity and Aembit both put a per-tool allow/deny decision in front of agent tool calls, so an MCP or Claude is blocked before it runs rather than reviewed after. GitHub's

vim ? help

agentic autofix for code scanning alerts hit GA. GitGuardian now checks vaulted and Kubernetes secrets live, so inventory entries carry a real verdict.

#### PROTECT

### **Block an agent's tool call before it executes, not after**

Both intercept at the point of invocation — Zenity via Anthropic's Inference Hooks across Claude Enterprise surfaces with no endpoint agent, Aembit with per-tool allow/block on MCP servers and fail-closed defaults. Agent governance has mostly meant logging what an assistant already did; this is the enforcement point that makes the policy actually binding, with denials visible and audited.

[Zenity](#) · [Aembit](#)

#### PROTECT

### **Ship a fix for a code scanning alert without writing the patch**

GitHub's agentic autofix for code scanning alerts is now GA, and GitLab 19.5 runs its Code Review Flow as a CI/CD job with custom instructions and self-hosted model support. The backlog of low-severity findings nobody has time to patch is the usual failure mode here; moving remediation into the pipeline is what makes clearing it cheap enough to actually happen.

[GitHub](#) · [GitLab](#)

#### IDENTIFY

### **Know which inventoried secrets and AI agents are actually live**

ggscout 0.34.1 validates vaulted and Kubernetes secrets from inside the network and puts live validity and permission verdicts on NHI inventory entries, so a stale vault record stops looking like an active credential. Runlayer's Detect now finds IDE-native AI agents and coding-assistant extensions that previously never appeared in any inventory at all.

[GitGuardian SaaS](#) · [Runlayer](#)

## Nmap

[Feedback](#)[v7.992](#)[CODE ↗](#)

Nmap - the Network Mapper. Github mirror of official SVN repository.

Nmap 7.992 adds a libfuzzer framework for internal function testing, NSE memory (400 MB) and CPU (5 s) guardrails, DHX2 AFP auth, and XML output for two NSE scripts.

↳ **GET THIS VERSION**

```
$ git clone --branch v7.992 https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout v7.992
```

- › Adds a libfuzzer-based fuzzing framework covering internal functions of Nmap, Ncat, and Nping, with 10 functions instrumented so far.
- › Limits NSE Lua interpreter to 400 MB memory and kills script threads that exceed 5 seconds of execution without yielding.
- › Adds `--script-timeout` and `--host-timeout` enforcement for previously uninterruptible NSE script loops in `rdp.lua`, `tn3270.lua`, and `ospf.lua`.
- › Replaces the cryptographically broken DHCAST128 UAM in `afp.lua` with DHX2, preventing credential decryption from network traffic.
- › Implements XML output for the `rsync-list-modules` NSE script.

+3 more

### COVERAGE

**NEW** detection [SAP HANA database service probe](#)

**UPDATED** supported platform [Bundled library upgrades](#)

[All of Nmap's coverage changes →](#)

WAS THIS USEFUL?



# StackHawk

Feedback



changelog · Oct 16, 2026

CHANGELOG ↗

RANK

StackHawk is an application security tool built for developers. Automated scanning for the entire dev pipeline combined with continuous documentation and remediation.

StackHawk shipped a Windows CLI installer, automatic OpenAPI base-path detection for scans, removal of the default per-rule alert cap, and two breaking changes to environment variable handling and authorization alpha rules.

↳ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

## 01 Windows CLI installer for HawkScan NEW 75

Adds a Windows installer script for the HawkScan CLI: `irm https://download.stackhawk.com/hawk/install.ps1 | iex` run in PowerShell, which installs to `%USERPROFILE%\hawk\bin`.

Install the HawkScan CLI on a Windows CI runner in a single PowerShell line.

```
$ irm https://download.stackhawk.com/hawk/install.ps1 | iex
```

– Runnable install command with exact install path changelog-20261016-0b95ee00

### ↳ BREAKING ON UPGRADE

- ! An unresolved `${VARIABLE}` placeholder in configuration now stops the scan with **Missing environment variable: VARIABLE** instead of silently substituting an empty string. Pipelines that intentionally left a variable unset must either set it or use the `${VARIABLE:}` (trailing-colon) form.
- ! Authorization alpha rules 40051, 40054, 40055, and 40056 have been removed; any configuration or workflow that relied on these single-identity rules will no longer produce results from them.

WAS THIS USEFUL?



Priority lanes, the lifecycle board, issue history, the Audit log page, and a regrouped

NOTES

RANK

ZeroPath is an AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives.

ZeroPath's October release reorganizes triage around exploit-evidence priority lanes and a drag-and-drop lifecycle board, adds an Audit log page, per-finding History and Dependencies tabs, a Repository risk overview, Organization SBOM export, and expanded issue search/API fields — alongside new permission requirements for re-investigation and tag management.

WHAT SHIPPED • 10 FEATURES most completely described first ? what's the number?

01 Issue search and reporting API field expansion IMPROVED

90

The issue search API now returns and filters by `lane` and `evidence`, sorts by `priority`, and returns linked tickets on issue list and search results. The API/MCP adds an `issues.vulnerabilityReport` field returning the per-finding report, and `presentation.evidence` to issue search when called with `presentation: true`, where `evidence.validated` covers agent-, Runtime Validation-, and person-confirmed findings (`validationProvenance: "legacy"` for findings confirmed before tracking); `knownVulnerableVersion: true` now marks version-match-only dependency findings without asserting `reachability: "reachable"`, a change from prior responses.

– Every API field named, directly callable, includes the response-shape change

Priority lanes, the lifecycle board, is...

02 Audit log page under Reporting NEW

85

A new `Audit log` page under `Reporting`, for organization admins, lets you filter events by actor, source, and level over up to 90 days, with per-event endpoint and argument detail, and `Export CSV` for up to 100,000 events.

Export the last 90 days of admin activity as CSV for a compliance audit — filter to a specific actor first to narrow scope.

In the ZeroPath app, go to Reporting > Audit log, set the date range to 90 days, filter by actor and level, then click `Export CSV` to download up to 100,000 events.

– Exact limits, filters and a worked export example make it fully actionable

Priority lanes, the lifecycle board, is...

### 03 Repository tag permission and API field changes

BREAKING

85

Creating, editing, and deleting repository tags now requires **Manage Teams** or **Edit SAST Configuration** permission (non-admins need a team granting **Manage Teams**), and tag create/update API requests that pass the **repositoryIds** field are now refused — callers must switch to explicit **include**, **exclude**, and **release** fields.

– Exact permissions and the replacement API fields are named

Priority lanes, the lifecycle board, is...

### 04 New report export options **NEW**

80

Adds **Organization SBOM** export bundling the latest SBOM for every viewable repository into one ZIP in **CycloneDX 1.6** or **SPDX 2.3** JSON with a manifest of included and skipped repositories, and a per-finding **Export report (.docx)** action under **More actions** on the full issue page.

Download the organisation-wide SBOM bundle (**CycloneDX** and **SPDX** for every repo) for an audit or supply-chain review.

📌 In the ZeroPath app, go to Reports > Generate, select **Organization SBOM**, choose CycloneDX 1.6 or SPDX 2.3 JSON, and click Generate. The resulting ZIP contains one SBOM per repository plus a manifest of included and skipped repositories.

– Exact formats plus a worked generation example Thin description, just a UI location and output format

Priority lanes, the lifecycle board, is...

### 05 Priority lanes and scan source filter on Issues page **NEW**

75

The Issues page now groups open findings into **Fix now**, **Needs decision**, and **Lower priority** lanes — **Fix now** requires High/Critical **CVSS** severity plus at least one exploit-evidence chip (Runtime Validation verdict, known exploited advisory, reachable dependency, Wiz internet exposure, or verified live secret) — and adds an **Issues scan source** menu offering **All**, **Full Scans**, and **Latest PR Scans** for code scans.

– Names lanes, exploit-evidence criteria and menu options but no exact nav path

Priority lanes, the lifecycle board, is...

### 06 Lifecycle board with accepted-risk expiry **NEW**

75

A new **Board view** lets findings move through **Needs review**, **Backlog**, **Reviewing**, and **Patching** stages via drag-and-drop, requiring a close-reason dialog and providing a 6-second Undo window; **Accepted risk** findings automatically reopen once their accepted-until date passes unless closed by other evidence.

– Mechanism and stages fully named, but only UI-level actionability

Priority lanes, the lifecycle board, is...

## 07 History and Dependencies tabs on findings NEW 75

Every finding now has a **History** tab showing status changes and detections in one timeline, naming the job (e.g. **ZeroPath risk acceptance expiry**) or ticket (e.g. **Jira · SEC-12**) behind each change, plus a **Dependencies** tab on the issue panel grouping advisories by package with an exploit probability chart and quick filters including **Known exploited** and **EPSS over 10%**.

– Concrete field examples given, but navigation only to a UI tab

Priority lanes, the lifecycle board, is...

## 08 API key management endpoints restricted to browser sessions NEW 70

New **Create**, **list**, and **delete** API key management endpoints are accessible only from a signed-in browser session; requests authenticated with an API key or **Authorization** header now receive **403** for these operations.

– Endpoints and exact status code named, access mechanism clear

Priority lanes, the lifecycle board, is...

## 09 Close Issues permission required for re-investigation 70

**BREAKING**

Requesting a re-investigation from the app, API, MCP, or Assistant now requires the **Close Issues** permission on the finding's repository; organizations using access control must explicitly grant it to members who re-investigate findings, or those requests will receive **403**.

– Names the exact permission and failure code, clear migration note

Priority lanes, the lifecycle board, is...

## 10 Repository risk overview on Repositories page NEW 65

A new **Repository risk overview** on the Repositories page lists repositories riskiest first, with columns for open findings, **Fix now** findings, exploit evidence, and **Protection** status (scheduled scans, PR scanning, PR comments).

– Named columns but only a page-level navigation hint

Priority lanes, the lifecycle board, is...

### ↳ **BREAKING ON UPGRADE**

- ! Requesting a re-investigation from the app, API, MCP, or Assistant now requires the **Close Issues** permission on the finding's repository; organizations using access control must

explicitly grant it to members who re-investigate findings, or those requests will receive `403` .

- ! Creating, editing, and deleting repository tags now requires `Manage Teams` or `Edit SAST Configuration` ; members who are not admins need a team that grants `Manage Teams` .
- ! Tag create and update API requests that pass the `repositoryIds` field are now refused; callers must switch to explicit `include` , `exclude` , and `release` fields.
- ! API key management ( `create` , `list` , `delete` ) now returns `403` for requests authenticated with an API key or an `Authorization` header; these operations require a signed-in browser session.
- ! A dependency finding confirmed only by a version match now carries `knownVulnerableVersion: true` and no longer reports `reachability: "reachable"` on that basis; callers relying on reachability for version-match findings will see a changed response.

---

WAS THIS USEFUL?



GitHub is a web-based platform for version control and collaborative software development using Git repositories.

GitHub brought agentic autofix for code scanning alerts to general availability, added multi-repo conversations to Copilot in Slack and Microsoft Teams, and let admins download Code Quality license usage as a CSV.

WHAT SHIPPED • 3 FEATURES most completely described first ? what's the number?

## 01 Agentic autofix for code scanning alerts reaches GA NEW

95

Agentic autofix is now generally available: it can be triggered via the [Update a Code Scanning Alert](#) REST API by setting `assignees` to `["copilot-swe-agent[bot]"]` on alerts on the default branch, which makes Copilot explore, fix, validate, and open a PR autonomously. Enterprise admins control access via a new [Agentic autofix](#) policy under Advanced Security policies, an [Assign to Copilot](#) button on individual alerts replaces the previous [Generate Fix](#) option for eligible customers, and bulk remediation of up to 25 alerts at once is supported from the security alerts list or a security campaign, consolidated into a single pull request.

Assign a specific code scanning alert to Copilot for autonomous remediation via the REST API, triggering file exploration, fix generation, CodeQL validation, and a draft PR.

```
$ curl -X PATCH \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  https://api.github.com/repos/<owner>/<repo>/code-scanning/alerts/<alert_number> \
  -d '{"assignees": ["copilot-swe-agent[bot]"]}'
```

Enable agentic autofix for an organization after the enterprise administrator has turned on the [Agentic autofix](#) policy, so Copilot can begin remediating code scanning alerts.

 In GitHub Enterprise Cloud, go to Enterprise Settings > Advanced Security policies > Agentic autofix and enable the policy. Then navigate to Organization Settings > Code Security and turn on Agentic autofix.

## 02 CSV export of Code Quality license usage NEW

80

Org and enterprise admins can download a CSV of GitHub Code Quality license usage from the **Licensing** section of the Billing and licensing page, available on GitHub Team and GitHub Enterprise Cloud. The report lists each active committer consuming a license, the repositories receiving their commits in the last 90 days, and whether that committer is unique to a single repository or organization.

Download a Code Quality license usage report to identify committers unique to a single repo where disabling Code Quality could free up licenses.

 In your organization or enterprise, go to Billing and licensing > Licensing > Code Quality card > Download CSV.

– Exact navigation path and report contents are specified

Download GitHub Code Quality license us...

## 03 Multi-repo conversations for Copilot in Slack and Teams NEW

68

Copilot in Slack and Microsoft Teams now supports querying across multiple repositories in a single conversation, enabling PR tracking, issue investigation, and activity summaries without switching conversations. **@GitHub** can now be mentioned anywhere in a message rather than only at the start, and Copilot uses shared Slack messages and threads as context so it can reference discussions without manual copy-paste, with smoother mid-conversation repository switching.

 [Watch video](#)

Get a cross-repo PR status summary without leaving Slack — useful at the start of a team standup.

 @GitHub Show me all my open pull requests across repositories and highlight which ones need my attention.

Investigate a bug that touches multiple services by asking Copilot in the context of an existing Slack thread about the incident.

 @GitHub Help me investigate an issue that spans our frontend and backend repositories.

– Clear example prompts given, but no config or API surface named

Multi-repo support and other improvemen...

## COVERAGE

**NEW**

supported platform [AI Scan language and framework coverage](#)

[All of GitHub's coverage changes →](#)

AI Scan for pull requests generally available (2026-10-15)

**BLOG** ↗

**STABLE**

AI Scan for PRs is GA: auto-runs on PRs, shows findings with autofix suggestions, requires Code Security or Advanced Security for private repos.

- › AI Scan runs automatically on pull requests when code scanning and the repository's effective AI Scan setting are enabled — no GitHub Copilot license required.
- › AI Scan findings appear alongside CodeQL results on the Conversation and Files changed tabs with an **AI** indicator, and can include an explanation of the security risk and a suggested autofix.
- › AI Scan for pull requests consumes Code Security AI Credits, with usage visible in the **AI usage** section of enterprise billing under the **Advanced Security** product; SKU-level budgets for Code Security AI Credits can also be set.
- › AI Scan is available for public repositories owned by organizations or personal accounts; private and internal organization-owned repositories require GitHub Code Security or GitHub Advanced Security.
- › AI Scan findings are advisory — they do not fail the code scanning check and cannot block pull request merges through ruleset requirements.

[+1 more](#)

WAS THIS USEFUL?



Connect your agents to your software lifecycle with the GitLab MCP server (2026-10-15)

BLOG

RANK

Teams and agents ship trusted software faster with GitLab, the intelligent orchestration platform with context and governance for the entire software lifecycle.

GitLab 19.5 ships agentic code review across the merge request lifecycle, headlined by a governed Code Review Flow that runs as a CI/CD job with custom instructions and self-hosted model support, alongside an upgraded Recommend Reviewers Flow and new Security Review, Fix CI/CD Pipeline, and Custom Flows.

WHAT SHIPPED • 3 FEATURES most completely described first ? what's the number?

01 Governed Code Review Flow via CI/CD NEW 86

GitLab Code Review Flow brings governed agentic code review to every merge request. It can be enabled automatically at the project, group, or instance level so every merge request is reviewed without manual assignment, with settings that cascade down the hierarchy, and it runs as a CI/CD job on GitLab-hosted or self-managed runners inside the same pipeline as the rest of delivery. Teams can supply custom review instructions — covering language conventions, security practices, and path-specific rules — by pointing the flow at a template project so standards travel with the work across a group, and can run it on self-hosted models to keep inference under customer control for air-gapped or regulated environments. A follow-up review rechecks prior findings from GitLab Duo Agent Platform, verifies fixes, and carries context forward from the earlier pass. It works across GitLab.com , GitLab Dedicated, and GitLab Self-Managed at a unit cost of \$0.25 per review.

- Names config levels, runner model, pricing; lacks exact flag or endpoint

Unblock agentic software delivery with ...

02 Recommend Reviewers Flow for merge requests IMPROVED 70

Recommend Reviewers Flow triggers automatically when a merge request is created, whether in draft or ready state, surfacing suggestions in a dedicated panel with one-click assignment and explaining the reasoning behind each recommendation so authors can approve or override with full context. It suggests the smallest set of human reviewers who can satisfy code-owner and approval rules, factoring in expertise, workload, availability, and timezone.

- Explains selection logic and UI flow but no config surface named

### 03 New agentic Flows: Security Review, Fix CI/CD Pipeline, Custom

NEW

48

GitLab 19.5 adds **Security Review Flow** to check merge request changes for business logic vulnerabilities before they merge, **Fix CI/CD Pipeline Flow** to diagnose and repair failed CI/CD jobs directly from the merge request, and **Custom Flows** to compose Recommend Reviewers Flow alongside other agentic workflows a team already runs.

– Names three new flows but gives no mechanism or setup path

Get every merge request in front of the...

#### COVERAGE

NEW integrations

GitLab MCP server GA — MCP-compatible client support, External Secrets Operator (Kubernetes), Terraform and OpenTofu data source, Vault-compatible API for OpenBao and Vault CLIs

NEW supported platforms

GitLab MCP server available on GitLab.com, GitLab Self-Managed, and GitLab Dedicated, GitLab Secrets Manager GA on GitLab.com and GitLab Self-Managed (Premium and Ultimate, 19.5).

All of GitLab's coverage changes →

Connect your agents to your software lifecycle with the GitLab MCP server (2026-10-15)

BLOG ↗

STABLE

GitLab MCP server is now GA in 19.5, adding tools for branches, semantic code search, pipeline inspection, artifact access, group browsing, wiki lookup, and full Duo Agent Platform session control — all enforced under existing GitLab permissions.

- › Adds **add\_branch** tool to the GitLab MCP server to start a feature branch from an agent.
- › Adds **get\_repository\_file** tool to read a file at any ref without extra API calls.
- › Adds **semantic\_search** tool to find code by meaning, answering 'where is this implemented?' without knowing the file or symbol name.
- › Adds **get\_pipeline** tool that returns a pipeline's status, jobs, downstream pipelines, and artifacts on request.
- › Adds **list\_pipelines** tool to surface patterns across recent pipeline runs.

+9 more

GitLab Secrets Manager: Govern secrets like your code (2026-10-15)

BLOG ↗

STABLE

GitLab Secrets Manager (GA, 19.5) scopes build-time secrets per job via the `secrets:` keyword in `.gitlab-ci.yml`, encrypts at rest with AES-256-GCM via OpenBao, streams full CRUD audit events with pipeline/job IDs, and inherits group-level secrets through existing project membership.

- › Declares secrets in `.gitlab-ci.yml` using the `secrets:` keyword, injecting values only into jobs that declare them and fall within scope (branch, protected-branch flag, `production/*` environment wildcards).
- › Streams every secret create, read, update, and delete event — including pipeline and job IDs — to audit event destinations and the GitLab audit trail.
- › Writes secret values to a temporary file inside the job and exposes only the file path, limiting exposure in subprocesses, crash dumps, and telemetry.
- › Encrypts secrets at rest with AES-256-GCM via OpenBao, keeping plaintext out of the GitLab application layer.
- › Issues short-lived tokens bound to each job's identity so jobs fetch values directly from OpenBao.

+5 more

WAS THIS USEFUL?



# Endor Labs

Feedback



Container Scanning + SCA = Better Together | Endor Labs (2026-10-16)

BLOG ↗

RANK

Endor Labs is an application security platform built for AI-generated and human-written code, helping teams prioritize and fix real risk across the SDLC.

Endor Labs shipped a new Agent Kit integration for SpaceXSI's Grok Build that brings AI SAST, SCA, and AURI remediation agents into agent coding sessions, paired with new Coding Agent Governance hooks and MCP/skill trust scoring to police agentic workflows.

↳ WHAT SHIPPED • 6 FEATURES most completely described first ? what's the number?

## 01 Agent Kit installation and deployment for Grok Build

NEW

78

The Endor Labs Agent Kit can be installed into a Grok Build IDE session via the marketplace or by typing `/add-plugin endorlabs` in the agent chat, bringing AI SAST, SCA, and remediation agents into the session. Agent Kit workflows can also run outside the IDE in CI or cloud runs through the Grok Build SDK.

Bootstrap Endor Labs security scanning inside an active Grok Build agent session without leaving the IDE.

```
$ /add-plugin endorlabs
```

– Exact install command given; SDK/CI path only briefly noted

Endor Labs + SpaceXSI: Securing every s...

## 02 Coding Agent Governance hooks for agent actions

NEW

65

Coding Agent Governance hooks record every MCP tool call, shell command, and file operation and enforce policies deterministically before each action runs, without relying on a model's judgment. The hooks also block confirmed malicious packages at install time during agent-driven dependency installation.

– Mechanism and scope described but no policy config surface named

Endor Labs + SpaceXSI: Securing every s...

## 03 AURI remediation agents for AI SAST and SCA

NEW

65

The Agent Kit includes specialized AURI remediation agents — an AI SAST Remediation agent and an SCA Remediation agent — that pull findings from scans, propose fixes, and open pull requests with developer approval.

– Named agents and PR workflow but no command or UI path shown

Endor Labs + SpaceXSI: Securing every s...

THINNER COVERAGE BELOW

04 **SCA Reachability analysis** IMPROVED 55

SCA Reachability analysis cuts findings by up to 92% by identifying which vulnerable dependencies the application actually calls.

– Concrete reduction figure but no usage steps given

Endor Labs + SpaceXSI: Securing every s...

05 **Endor Scores for MCP servers and skills** NEW 50

Endor Scores rate each MCP server and skill for trust, mapped to the OWASP Top 10 for LLM Applications, giving security teams a basis to approve or block agents and skills.

– Names the scoring framework but not the approval workflow

Endor Labs + SpaceXSI: Securing every s...

06 **AI SAST data-flow tracing and secret detection** IMPROVED 50

AI SAST traces data flow from untrusted input to vulnerable calls in first-party code and flags exposed secrets before they are committed.

– Describes mechanism but no scan command or scope given

Endor Labs + SpaceXSI: Securing every s...

COVERAGE

NEW vulnerability check

Container image OS and application package vulnerability scanning

NEW integrations

Runtime CNAPP integration for container tracing and admission control, Pipeline Discovery — CI/CD tool and SCM visibility

NEW detection

Repository Security Posture Management (RSPM) out-of-the-box policies

NEW compliance framework

SLSA, NIST SSDF, SOC2, PCI alignment via RSPM and artifact signing

All of Endor Labs's coverage changes →

STABLE

Endor Labs adds container image signing at build time, signature verification at deploy time, and unified cross-dependency PR comments for vulnerability findings.

- › Adds container image signing at build time and signature verification at deploy time, with scan results consumable by CNAPP and admission controllers.
- › Includes container vulnerability findings in pull-request comments alongside application dependency findings, so developers see all impacted packages in one place.
- › Enables tracing of a production container finding back to its source repo, release artifact, and build via integration with runtime CNAPP tools.

Introducing CI/CD Security with Endor Labs | Endor Labs (2026-10-16)

BLOG ›

LATEST

STABLE

Endor Labs expands its SSCS platform with CI/CD security: pipeline discovery, repository posture management, and artifact signing with provenance.

- › Adds Pipeline Discovery to map every tool touching code across repositories and CI/CD workflows, surfacing shadow engineering, security coverage gaps, and policy violations.
- › Adds Build Integrity Verification via artifact signing, attaching provenance information (source repository, branch, commit, repository owner) to signed packages and containers.

WAS THIS USEFUL?



# Cloudsmith

Feedback



Proxy and cache packages from your ActiveState Curated Catalog (2026-10-15)

**BLOG** »

Cloudsmith is a fully-managed artifact management, software supply chain security, and software distribution platform.

## COVERAGE

**NEW**

integration [ActiveState Curated Catalog upstream](#)

[All of Cloudsmith's coverage changes →](#)

WAS THIS USEFUL?



JFrog software supply chain solutions provide an end-to-end pipeline to control your binaries from build to production. Power your software updates to the edge.

JFrog introduced the Business Value Center, a real-time ROI and supply chain coverage dashboard built into MyJFrog for SaaS customers.

← WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

## 01 Business Value Center dashboard in MyJFrog NEW 68

A new real-time dashboard inside **MyJFrog** quantifies platform ROI, developer hours reclaimed, and storage saved via deduplication, alongside the percentage of repositories actively scanned by JFrog Xray. It also surfaces an Optimization panel showing what share of the software supply chain is scanned for vulnerabilities and what fraction of incoming open-source packages are blocked by JFrog Curation policies before reaching developers.

Access the Business Value Center to pull real-time ROI figures before a budget renewal or executive review.

📌 Sign in to MyJFrog > select **Business Value Center** from the dashboard navigation > review the ROI panel (developer hours reclaimed, cost avoided, storage saved) and the Optimization panel (% of repos scanned, % of packages blocked by Curation policies).

– Names specific metrics and UI path but no API or config surface

JFrog Business Value Center: Celebrate ...

WAS THIS USEFUL?



# infisical

Feedback



v0.166.8

NOTES ↗

RANK

Infisical is the open-source platform for securing secrets, certificates, access, and AI Agents.

Infisical v0.166.8 adds Portainer as a secret sync target, strengthens webhook alerting with authentication and custom headers plus a new audit-log-stream drop alert, and fixes tag visibility during secret query evaluation.

↳ WHAT SHIPPED · 3 FEATURES most completely described first ? what's the number?

## 01 Portainer secret sync target NEW 35

Infisical adds Portainer as a supported secret sync app connection target, allowing secrets to be synced to Portainer environments.

– Names the integration but no setup steps or scope given v0.166.8

## 02 Webhook and audit log alerting improvements IMPROVED 35

Webhook alert channels now support authentication and custom headers, and a new alert fires when an audit log stream starts dropping logs.

– Describes what changed but no config keys or UI path given v0.166.8

## 03 Tag visibility in secret query evaluation IMPROVED 20

Secret query evaluation now returns all tags, fixing a gap where tags were previously omitted.

– Bare bug-fix description with no mechanism or surface named v0.166.8

WAS THIS USEFUL?



GitGuardian SaaS is a cloud-hosted security platform that detects exposed secrets, API keys, and credentials in source code repositories and developer workflows.

GitGuardian shipped in-network validity and permissions checking for vaulted and Kubernetes secrets via ggscout 0.34.1, surfacing live verdicts in the [NHI](#) inventory, alongside a new saved view in Endpoint Protection, fuller API permission docs, and priority queuing for Jira Cloud historical scans.

←→ [WHAT SHIPPED](#) · [4 FEATURES](#) most completely described first [? what's the number?](#)

## 01 In-network validity checks for vaulted secrets in NHI inventory

NEW

95

ggscout 0.34.1 checks validity and permissions of secrets held in vaults or Kubernetes from inside the network. It adds a [Validity](#) column and filter (Valid, Invalid, Unknown) in [NHI Governance](#) → Identities, surfaces token permission scopes (e.g. [repo](#), [admin:org](#) on GitHub; [api](#) on GitLab) alongside each identity, and reports credential owner, expiry date, and provider metadata sourced directly from the provider during each ggscout run. Validity verdicts from vaulted or Kubernetes secrets are merged with any matching secret incident, always keeping the most recent verdict, and the [\[check\]](#) [enabled = false](#) config key lets a deployment opt out of these checks.

Disable in-network validity checks for a specific ggscout deployment that cannot reach external providers.

```
ini

[check]
enabled = false
```

After upgrading to ggscout 0.34.1, run a fetch-and-send to populate validity verdicts, then filter the NHI inventory to immediately see all invalid credentials still sitting in vaults or Kubernetes clusters.

[In NHI Governance](#) → Identities, open the Validity filter and select [Invalid](#) to list revoked credentials across vaults, CI variables, and Kubernetes clusters.

– Names config key, exact filter values, scopes, and navigation path

Check validity and permissions of your ...

## 02 Priority queuing for Jira Cloud historical scans IMPROVED

75

Adds on-demand priority queuing for Jira Cloud historical scans launched from the dashboard or public API so they start ahead of recurrent and system scans, and counts public API historical scan launches made with a service account token as on-demand on AWS S3 and version control sources.

– Explains queuing behaviour and named sources but no config surface

Check validity and permissions of your ...

THINNER COVERAGE BELOW

## 03 Built-in 'Valid' saved view in Endpoint Protection NEW

45

Adds a built-in **Valid** saved view on the Endpoint Protection discovered secrets page, available to all users.

– Names the exact UI surface but no mechanism

Check validity and permissions of your ...

## 04 Full API token permissions documentation IMPROVED

40

Documents all API token permissions in the public API documentation, including eleven previously granted but undocumented permissions.

– States scope and a count but no specifics of the permissions

Check validity and permissions of your ...

WAS THIS USEFUL?



# Invicti

Feedback



Release 20261015 (2026-10-15)

BLOG

RANK

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

Invicti's October 2026 release adds bulk vulnerability retesting and API key authentication for CI/CD Source Scans, alongside improvements to API catalog visibility, login sequence troubleshooting, permission scoping, and SCIM user management.

WHAT SHIPPED • 7 FEATURES most completely described first ? what's the number?

## 01 Bulk retest of scan vulnerabilities NEW 70

Adds a **Retest found vulnerabilities** option in the **Scan again** dropdown on the scan summary, letting a user retest every vulnerability found by a scan in one action instead of retesting each individually.

After deploying a fix, confirm all vulnerabilities from a scan are resolved without retesting each one individually.

In the scan summary, open the **Scan again** dropdown and select **Retest found vulnerabilities**.

— Names exact UI control and action, but no mechanism detail beyond that

Release 20261015 (2026-10-15)

## 02 API key authentication for Source Scan NEW 70

Adds API key authentication for Source Scan in CI/CD pipelines, with expiry date support configured on the **Profile API key** page, so pipelines no longer depend on short-lived tokens.

Authenticate a long-running pipeline Source Scan with an API key rather than a short-lived token, and set an expiry date for the key.

Navigate to **Profile API key**, generate an API key for Source Scan, set the desired expiry date, and configure your CI/CD pipeline to pass the key as the authentication credential.

— Exact page named and steps to configure provided

Release 20261015 (2026-10-15)

## 03 Login sequence recording and playback diagnostics IMPROVED 60

Adds HAR file export from login sequence recording sessions, with sensitive headers masked, to assist Support troubleshooting, and adds specific failure messages when login sequence playback, session recovery, or automatic login fails, replacing generic errors.

– Names two concrete additions but no exact UI steps given Release 20261015 (2026-10-15)

THINNER COVERAGE BELOW

04 **API catalog target display and filtering** IMPROVED 45

Improves the API catalog by adding target name and URL display with filtering by those fields, and by showing how each API was discovered via a distinct icon per API source.

– Describes what changed but no exact navigation given Release 20261015 (2026-10-15)

05 **Scoped API specification update permission** IMPROVED 45

Scopes the API specification update permission to asset update permissions, allowing the ability to be granted without requiring full API Security access.

– Names the permission change but no UI path to grant it Release 20261015 (2026-10-15)

06 **Bulk removal of SCIM-provisioned users** NEW 35

Adds bulk removal of SCIM-provisioned users from Invicti Platform in a single operation.

– States the action but no steps or limits given Release 20261015 (2026-10-15)

07 **Fail scans on authentication failure** NEW 30

Adds a setting to fail a scan when authentication problems occur, preventing unauthenticated results from being silently accepted.

– No config key or UI path named for the setting Release 20261015 (2026-10-15)

↳ **ALSO FROM THESE RELEASES**

.  
. .  
.

WAS THIS USEFUL?  

# Fastly

[Feedback](#)

CLI v16.2.0 (2026-10-15)

[BLOG ↗](#)[RANK](#)

Fastly is an edge cloud platform offering content delivery network, load balancing, edge computing, and web application security services for developers and enterprises.

Fastly's CLI v16.2.0 adds a new routing-config subcommand, GCS logging public-key support, and a shorter help flag.

↳ WHAT SHIPPED · 3 FEATURES most completely described first ? what's the number?

## 01 GCS logging public-key flag NEW

70

Adds a `--public-key` flag to GCS logging commands ( `service/logging` ), allowing a public key to be attached when creating or updating a GCS logging endpoint.

Attach a GCS public key when creating or updating a GCS logging endpoint for a service.

```
$ fastly logging gcs create --service-id $FASTLY_SERVICE_ID --  
version active --name my-gcs-endpoint --bucket my-bucket --  
public-key "$(cat gcs-public-key.pem)"
```

– Named flag with a runnable example command CLI v16.2.0 (2026-10-15)

THINNER COVERAGE BELOW

## 02 Short help flag alias NEW

40

Adds `-h` as a short alias for `--help` across all commands.

– Simple named flag, minimal further detail needed CLI v16.2.0 (2026-10-15)

## 03 Routing config subcommand in CLI NEW

35

Adds CLI support for routing configs via the new `routing-config` subcommand.

– Names subcommand but no usage example or further mechanism given

CLI v16.2.0 (2026-10-15)

WAS THIS USEFUL?



# Aviatrix CoPilot

Feedback



New and Enhanced Features in 4.38.0

NOTES ↗

RANK

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

CoPilot 4.38.0 adds a configurable login banner and route approval email alerts, alongside enhancements to the gateway upgrade workflow, gateway grid/routing, and the Distributed Cloud Firewall.

←→ WHAT SHIPPED • 5 FEATURES most completely described first ? what's the number?

## 01 Route approval email alerts NEW 35

CoPilot adds route approval email alerts that notify administrators when routes require approval.

– Clear trigger and behavior, no setup steps given New and Enhanced Features in 4.38.0

## 02 Login banner configuration NEW 33

CoPilot now supports login banner configuration, letting administrators customize the authentication screen.

– Names the capability but no config key or UI path given

New and Enhanced Features in 4.38.0

## 03 Distributed Cloud Firewall enhancements IMPROVED 21

Delivers Distributed Cloud Firewall enhancements.

– Names the feature area but no concrete change described

New and Enhanced Features in 4.38.0

## 04 Gateway upgrade workflow enhancements IMPROVED 18

Delivers upgrade experience enhancements for the gateway upgrade workflow.

– Vague mention with no detail on what changed New and Enhanced Features in 4.38.0

## 05 Gateway grid and routing enhancements IMPROVED 18

Delivers gateway grid and routing enhancements.

– No specifics on which grid or routing behaviors changed

New and Enhanced Features in 4.38.0

↳ ALSO FROM THESE RELEASES

.  
.  
.

---

WAS THIS USEFUL?



---

**DETECTION & RESPONSE**

# isMalicious

Feedback



changelog · Oct 16, 2026

CHANGELOG ↗

RANK

IsMalicious is a threat intelligence platform that checks IPs, domains, URLs, emails, phones and crypto wallets against indexed threat indicators.

isMalicious added CVE intelligence pages with correlated risk scoring, new STIX/TAXII 2.1 feeds, DNS history lookups, expanded ransomware victim tracking, webhook retry logging, and a searchable security glossary.

↳ WHAT SHIPPED · 3 FEATURES most completely described first ? what's the number?

## 01 Threat intelligence glossary NEW 50

Adds a searchable glossary of 50+ terms spanning threat intelligence, malware, network, vulnerability, and API security, including a new AI agent security category, plus long-form explanations and worked examples for 16 top terms and 14 additional new terms.

– Scope and counts given but no page path or search mechanism [changelog-20261016-9e5d602c](#)

## 02 Ransomware victim tracking via RansomLook and Ransomfeed IMPROVED 48

Expands ransomware victim post coverage via RansomLook and Ransomfeed, with each entry carrying the ransomware group, publication time, and correlated indicators.

– Names sources and fields but gives no lookup path or endpoint

[changelog-20261016-9e5d602c](#)

## 03 DNS history NEW 21

Adds DNS history lookups.

– Named only in the summary with no mechanism or scope detail [changelog-20261016-9e5d602c](#)

WAS THIS USEFUL?



# Databricks

[Feedback](#)[READ METADATA privilege support for compliance security profile standards \(2026-10-16\)](#)[BLOG ↗](#)[RANK](#)

Databricks is a cloud-based data lakehouse platform for data engineering, analytics, machine learning, and AI workloads, built on Apache Spark.

Databricks promoted the [ADBC](#) driver for its Power BI connector to general availability, making it the default over [ODBC](#) for new connections.

←→ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

## 01 ADBC driver for Power BI connector NEW 55

The Arrow Database Connectivity (ADBC) driver for the Databricks connector for Power BI is now generally available and is set as the default driver for new connections, replacing ODBC, in both Power BI Desktop and the Power BI service.

– Names driver and default behaviour but no config/migration steps

ADBC driver for Power BI is now general...

### COVERAGE

UPDATED compliance framework

[READ METADATA privilege support for compliance security profile standards](#)

[All of Databricks's coverage changes →](#)

[READ METADATA privilege support for compliance security profile standards \(2026-10-16\)](#)[BLOG ↗](#)LATESTSTABLE

The READ METADATA privilege is now enabled by default for workspaces with the compliance security profile active.

- › Enables the READ METADATA privilege by default for workspaces with the compliance security profile enabled, covering all supported compliance standards.

WAS THIS USEFUL?



# Wazuh

[Feedback](#)[v4.14.9](#)[NOTES](#)[RANK](#)

Wazuh - The Open Source Security Platform. Unified XDR and SIEM protection for endpoints and cloud workloads.

Wazuh v4.14.9 increases the active response command capacity for wazuh-execd, extends FIM eBPF whodata support to older kernels, and tightens API authentication-context handling with a configurable payload limit and safer access logging.

→ WHAT SHIPPED • 4 FEATURES most completely described first ? what's the number?

## 01 Configurable API auth-context payload size IMPROVED

80

Adds the `auth_context_max_payload_size` option to configure the API `run_as` authentication-context payload size limit, raised from 8 KB to 64 KB by default, improving support for AD/LDAP/SSO logins that carry large group-membership contexts.

– Names exact config key and before/after size limits v4.14.9

## 02 Active response command capacity raised to 1024 IMPROVED

60

Raises from 64 to 1024 the number of active response commands that `wazuh-execd` can load from `etc/shared/ar.conf`.

– Names component, config file, and exact numbers v4.14.9

## 03 API access log masks run\_as context IMPROVED

60

Restricts the API access log from writing the `run_as` authorization context verbatim unless debug logging is enabled; `hash_auth_context` still identifies it.

– Names field and condition, no example of log output v4.14.9

THINNER COVERAGE BELOW

## 04 FIM eBPF whodata on older kernels IMPROVED

40

Allows FIM eBPF whodata on capable kernels older than 5.8, such as RHEL 8.10.

– Names kernel version and OS but no further mechanism v4.14.9

WAS THIS USEFUL?



# Google SecOps

Feedback



October 12, 2026 (2026-10-12)

[BLOG ↗](#)

Google secops is a comprehensive documentation, guides, and resources for {% dynamic print site\_values.cloud\_name %} products and services.

## ↳➤ BREAKING ON UPGRADE

! Custom rules referencing `MANDIANT_ACTIVE_BREACH_IOC` , `MANDIANT_FUSION_IOC` , or `OPEN_SOURCE_INTEL_IOC` must be updated to use `GTI_IOC` .

## COVERAGE

**UPDATED** detection

[Applied Threat Intelligence curated rules migrated to GTI\\_IOC feed](#)

[All of Google SecOps's coverage changes →](#)

WAS THIS USEFUL?



# Trend Vision One

Feedback



Zero Trust Secure Access Private Access Connector Version 3.11.0.2190 (2026-10-16)

BLOG ↗

RANK

Trend Vision One is Trend Micro's unified cybersecurity platform that consolidates threat detection, response, and risk management across endpoints, cloud, email, and networks.

Trend Vision One's window centers on email security, with a new option to request attachment passwords directly from recipients, alongside a minor connector update adding image hardening and TCP MSS clamping.

←→ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

## 01 Recipient password prompts for protected attachments

NEW

70

In Email and Collaboration Security → Cloud Email and Collaboration Protection, a new policy action requests the password for a password-protected attachment directly from the recipient when the password is absent from the email, enabling scanning of attachments that would otherwise be skipped as unscannable, without administrator intervention.

Enable end-user password prompting so password-protected attachments are scanned rather than logged as unscannable.

📌 In the console, go to Email and Collaboration Security > Cloud Email and Collaboration Protection, locate the policy action for password-protected attachments, and select the option to request the attachment password from the recipient.

– Names exact UI path and mechanism, includes usage steps

Password-protected attachment support w...

THINNER COVERAGE BELOW

## 02 Private Access Connector hardening update

IMPROVED

35

Private Access Connector 3.11.0.2190 adds TCP MSS clamping and applies additional image hardening settings recommended by CIS benchmarks.

– Names version and two changes but no further mechanism

Zero Trust Secure Access Private Access...

WAS THIS USEFUL?



◆ SOAR

## Tines

Feedback



Keyless cloud credentials with workload identity federation (2026-10-15)

BLOG ↗

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 adds keyless cloud credential authentication via workload identity federation for AWS, Google Cloud, Microsoft Entra, and Oracle Cloud.

- › Supports workload identity federation for keyless authentication with AWS, Google Cloud, Microsoft Entra, and Oracle Cloud credentials.

### COVERAGE

NEW

integration Workload identity federation support for cloud credential providers

All of Tines's coverage changes →

WAS THIS USEFUL?



# Tenzir

Feedback



Tenzir Node v6.20.0-rc.3: A new data model for faster pipelines (2026-10-15)

BLOG ↗

RANK

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

Tenzir v6.20.0-rc.3 introduces a new data model with typed AI decisions via a new `ai_decide` operator, secret-backed pipeline expressions, and a broad set of pushdown and integration improvements spanning ClickHouse, SentinelOne, MySQL, SQL Server, Windows Event Forwarding and Google Cloud Monitoring.

↳ WHAT SHIPPED • 9 FEATURES most completely described first ? what's the number?

## 01 ClickHouse integration: map columns, catch-all tables, append default

80

BREAKING

Improves ClickHouse integration with `from_clickhouse` `Map columns` support via `map columns in from_clickhouse`, reading from ClickHouse `catch-all tables`, and `Append as default mode for to_clickhouse` — this new default write mode changes behaviour for existing pipelines that relied on the previous default.

– Names operators and flags, includes before/after of default change

Tenzir Node v6.20.0-rc.3: A new data mo...

## 02 Typed AI decisions with `ai_decide` operator

NEW

70

Adds `ai_decide` operator with typed decisions and dedicated `Decisions` endpoints, so pipelines can branch on structured AI-evaluated results.

– Names operator and endpoints but no usage example

Tenzir Node v6.20.0-rc.3: A new data mo...

## 03 SentinelOne integration: Data Lake auth, pushdown, console API

IMPROVED

65

Improves the SentinelOne integration with `SentinelOne Data Lake API keys` authentication, `Filter and time-window pushdown` for SentinelOne queries, and `SentinelOne queries through the console API` as an alternative query path.

– Three named improvements, no exact config keys

Tenzir Node v6.20.0-rc.3: A new data mo...

## 04 Windows Event Forwarding ingestion and Google Cloud Monitoring sink

NEW

60

Adds **Windows Event Forwarding** ingestion without requiring a Windows collector and **Write metrics to Google Cloud Monitoring** as a new metrics sink.

– Names two surfaces but no config detail Tenzir Node v6.20.0-rc.3: A new data mo...

## 05 Pushdown enhancements across six backends and MySQL/SQL Server

NEW

60

Adds pushdown support across backends, including **Summaries computed in DuckDB**, **ClickHouse**, **MySQL**, **SQL Server**, **SentinelOne**, and through source renames/projections, plus **Filter, projection, and limit pushdown** for MySQL and SQL Server data sources.

– Names all six backends but no mechanism depth Names sources and pushdown types, no example

Tenzir Node v6.20.0-rc.3: A new data mo...

THINNER COVERAGE BELOW

## 06 Secrets as pipeline function arguments

NEW

55

Adds secrets as function arguments and **Secret keys for hmac**, enabling secret-backed HMAC signing inside pipeline expressions.

– Describes mechanism but no syntax shown Tenzir Node v6.20.0-rc.3: A new data mo...

## 07 New built-in functions: Luhn, CRC-32, OCSF, encryption

NEW

55

Adds **Luhn check digits**, **CRC-32 checksums**, **OCSF functions**, and **Encryption functions** as new built-in functions for pipeline expressions.

– Names functions but no behaviour detail per function

Tenzir Node v6.20.0-rc.3: A new data mo...

## 08 Named in-memory DuckDB databases

NEW

50

Adds **Named in-memory DuckDB databases** with read/write support via the DuckDB operators.

– Names capability, lacks operator syntax example

Tenzir Node v6.20.0-rc.3: A new data mo...

## 09 Parallel load balancing across pipeline branches

NEW

35

Adds **Parallel load balancing** across pipeline branches.

– Bare name, no mechanism or configuration detail

### ↳ **BREAKING ON UPGRADE**

! **Append** is now the default mode for **to\_clickhouse**, which changes the write behaviour for existing pipelines that previously relied on a different default.

---

WAS THIS USEFUL?



---

## EXPOSURE & VULNERABILITY MANAGEMENT

## Iru

Feedback



Automatically rotate FileVault recovery keys after viewing

NOTES ↗

RANK

Iru is a device-centric platform that unifies security, IT, and compliance around devices and everything they touch.

Iru shipped auto-rotation of viewed FileVault recovery keys, AI-assisted Library Item creation for BitLocker and Android Restrictions, and a separate target OS version setting for Automated Device Enrollment.


←→ WHAT SHIPPED • 3 FEATURES most completely described first ? what's the number?

### 01 AI-assisted Library Item creation for BitLocker and Android Restrictions

NEW

72

Iru AI can generate Library Items from a plain-language description for two policy types: BitLocker disk encryption policies on Windows, and Android Restrictions for managed devices. Admins use **Create with AI** in Library Items, pick the policy type, and describe the desired policy (e.g. requiring TPM-only encryption with a 256-bit AES key, or disabling the camera and blocking app installs from unknown sources).



Generate a BitLocker encryption policy Library Item without manually configuring settings — just describe the policy in plain language.

📌 In Iru, navigate to Library Items, select **Create with AI**, choose **BitLocker**, then describe your disk encryption policy (e.g. 'Require TPM-only encryption on all Windows endpoints and enforce a 256-bit AES key') and confirm.

Generate an Android Restrictions Library Item for managed devices by describing the restrictions you need in plain language.

📌 In Iru, navigate to Library Items, select **Create with AI**, choose **Android Restrictions**, then describe the restrictions (e.g. 'Disable camera, block app installs from unknown sources, and enforce screen lock after 30 seconds') and confirm.


– Clear UI steps and concrete example prompts for both policy types

Iru AI assistant for BitLocker and Andr...

## 02 Auto-rotate FileVault recovery keys after viewing NEW

65

Iru now auto-rotates FileVault recovery keys after an administrator views the current key, configurable via the **Automatically rotate keys after viewing** setting in the FileVault Library Item, with adjustable rotation timing. Key view events continue to be recorded in the device activity timeline.



Ensure a viewed FileVault recovery key is never reusable — enable auto-rotation so a new key is escrowed immediately after any admin view.

 In the Iru console, go to the FileVault Library Item > turn on **Automatically rotate keys after viewing** > set the desired rotation timing > save the Library Item.


– Names the exact setting and config path, lacks rotation timing details

Automatically rotate FileVault recovery...

## 03 Target OS version in Automated Device Enrollment NEW

62

Automated Device Enrollment gains an **Update to** setting letting admins define a separate target OS version that devices upgrade to before enrolling, independent of the existing minimum required OS version floor. Supported for macOS, iOS, and more.



– Names the exact setting but gives no navigation steps

Target OS version for Automated Device ...

### COVERAGE

NEW supported platform [Platform SSO OpenID web sign-in for macOS 27](#)

**UPDATED** integration

## Android Enterprise integration — managed Google domain upgrade

All of Iru's coverage changes →

Platform SSO OpenID and web login settings for macOS 27

**NOTES** ↗

**STABLE**

Single Sign-On Library Item gains OpenID web sign-in on macOS 27, with new settings to sync web login passwords and configure a URL allow list.

- › Adds OpenID sign-in support to the Single Sign-On Library Item for macOS 27 and later, enabling identity provider web sign-in pages (including MFA) at the login window, Lock Screen, and FileVault unlock.
- › New setting to sync the web login password to the local macOS account.
- › New setting to configure a web login URL allow list.

Upgrade your Android Enterprise integration to a managed Google domain

**NOTES** ↗

**LATEST**

**STABLE**

WAS THIS USEFUL?



# Axonius

[Feedback](#)[Axonius Release Notes 10.0.0 \(2026-10-11\)](#)[BLOG ↗](#)[RANK](#)

Axonius is a cybersecurity asset management platform that aggregates device and user data from many security and IT tools to find coverage gaps and enforce security policies.

Axonius 10.0.0 introduces a full AI suite — AI Companion, AI Workforce, and an MCP Server — alongside new correlation-investigation tooling, AI-specific attribute fields for software and SaaS assets, and several security-findings and exception-management additions.

←→ WHAT SHIPPED • 10 FEATURES most completely described first ? what's the number?

## 01 Status Resolver for stale security findings NEW 75

Adds a **Status Resolver** option in **Global Settings → Security Findings Data → Status Resolver** to automatically remove security findings closed for more than a configurable number of days (default 14 days for new customers, off for existing). Example: enable the auto-remove option and set the threshold to 30 days to keep the findings inventory from accumulating stale closed items.

Prevent stale, long-closed findings from cluttering your Security Findings inventory by enabling automatic removal after 30 days.

📌 In the Axonius console, go to Global Settings → Security Findings Data → Status Resolver, enable the auto-remove option, and set the threshold to 30 days.

– Exact config path, default value, and a runnable example

[Axonius Release Notes 10.0.0 \(2026-10-1...](#)

## 02 Explain Correlation and Non-Correlation modules NEW 70

Adds Explain Correlation and Explain Non-Correlation modules to investigate why device records were merged or kept separate, with AI-assisted deep-dive for Axonius AI customers. On the Devices page, selecting two device records and choosing Explain Non-Correlation shows the specific field discrepancies that prevented the merge.

Audit why two device records were not merged to resolve an unexpected duplicate before opening a support ticket.

📌 On the Devices page, select two device records, then choose Explain Non-Correlation to view the specific field discrepancies that prevented the merge.

– Includes mechanism and exact UI steps from the example

**03 Preferred Device Type field** **IMPROVED**

65

Adds a **Preferred: Device Type** preferred field, calculated post-correlation, usable across Query Wizard, reports, dashboards, and Enforcement Sets — replacing the legacy Device Type identifier field.

– Names field, usage surfaces, and replaced legacy field

Axonius Release Notes 10.0.0 (2026-10-1...

**04 Customizable email templates for Exception Management** **NEW**

65

Adds customizable email templates for Exception Management, supporting custom subject, intro text, and rich-text body for seven exception email types: approval request, approval reminder, request approved, request denied, request update, expiration warning, and expired.

– Enumerates all seven email types and customizable elements

Axonius Release Notes 10.0.0 (2026-10-1...

**05 Axonius AI Companion for natural-language queries** **NEW**

60

Introduces Axonius AI Companion for querying asset data in natural language, translating plain-English questions into queries, charts, dashboards, and actions with a visible correlation path for analyst verification.

– Explains mechanism and outputs but no exact command or path

Axonius Release Notes 10.0.0 (2026-10-1...

**06 AI attribute fields for software and SaaS assets** **NEW**

60

Adds seven AI attribute fields — **AI: Deployment**, **AI: Model Location**, **AI: Autonomy Level**, **AI: Integrations**, **AI: RAG Enabling**, **AI: RBAC Supported**, **AI: Vendor Training** — automatically populated on Software, installed software, and SaaS Application assets.

– Names all seven fields and asset types, no usage steps

Axonius Release Notes 10.0.0 (2026-10-1...

- 07 **Aggregated security findings status counts** NEW  60
- Adds three new count fields to Aggregated Security Findings — [Open Security Findings](#) , [Closed Security Findings](#) , and [Risk Accepted Security Findings](#) — as a complete, non-overlapping partition of Axonius Status values.

— Names fields and their logical structure, no usage path given

Axonius Release Notes 10.0.0 (2026-10-1...

- 08 **Managed SaaS auto-approval setting** NEW  60
- Adds an [Auto Approve Managed SaaS Applications](#) setting in [Settings → Enrichment](#) to control whether managed SaaS applications receive automatic approval status.

— Names setting and exact location, behavior kept brief

Axonius Release Notes 10.0.0 (2026-10-1...

THINNER COVERAGE BELOW

- 09 **Axonius MCP Server for external AI tools** NEW  50
- Introduces Axonius MCP Server, exposing the full Axonius environment to external AI tools, agents, IDE assistants, and low-code platforms via the Model Context Protocol.

— Names protocol and consumer types but no setup steps

Axonius Release Notes 10.0.0 (2026-10-1...

- 10 **Axonius AI Workforce automation engine** NEW  40
- Introduces Axonius AI Workforce, a platform-wide automation engine for end-to-end ticket creation, policy enforcement, and case updates driven by asset context.

— Describes purpose only, no concrete mechanism or entry point

Axonius Release Notes 10.0.0 (2026-10-1...

↳ **ALSO FROM THESE RELEASES**

.

WAS THIS USEFUL?



**IDENTITY, ACCESS & DATA PROTECTION**

# Varonis

[Feedback](#)[changelog · Oct 16, 2026](#)[CHANGELOG ↗](#)[RANK](#)

Varonis is the leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

Varonis this window added scope controls for AI classification, automated Databricks column masking from policy, configurable collector scan ratios, Snowflake database sizing, and support for migrating Salesforce integrations to Salesforce's External Client App framework.

↳ WHAT SHIPPED · 3 FEATURES most completely described first ? what's the number?

## 01 Collector scan ratio configuration IMPROVED 80

Collector settings now expose a configurable initial-vs-incremental scan ratio (e.g. 80/20 initial vs. incremental) instead of the default 50/50 balanced allocation, letting admins prioritize faster initial scan completion in newly onboarded environments with large data volumes.

Prioritize faster initial scan completion in a newly onboarded environment with large data volumes by shifting collector capacity away from incremental scans.

📌 In Varonis collector settings, adjust the scan ratio to allocate more capacity to initial discovery (e.g. 80/20 initial vs. incremental) rather than the default 50/50 balanced allocation.

– Names exact config behaviour and default/override ratios [changelog-20261016-e89b2684](#)

THINNER COVERAGE BELOW

## 02 Salesforce External Client App migration IMPROVED 40

Varonis supports migrating existing Salesforce integrations to Salesforce's External Client App framework while preserving existing configurations.

– Describes what but no steps or mechanism [changelog-20261016-e89b2684](#)

## 03 Snowflake database sizing NEW 15

Varonis added Snowflake DB sizing capability, per the changelog summary; no further detail provided.

– Only named in summary, no detail or example given [changelog-20261016-e89b2684](#)

WAS THIS USEFUL?



Use Immuta to provision governed access to data.

Immuta shipped a breaking change making tag-driven domain update endpoints asynchronous, and tightened authentication security by dropping SHA-1 support in SAML and unsigned JWTs in OpenID Connect.

←→ **WHAT SHIPPED** · **2 FEATURES** most completely described first ? what's the number?

## 01 Tag-driven domain updates now asynchronous

**BREAKING**

65

Tag-driven domain update API endpoints now return **202 Accepted** instead of **200 OK**. Callers that previously assumed the update was complete when the request returned must be updated to handle the async response.

– Names exact status code change but no endpoint path given. [snapshot-20261016](#)

## 02 Insecure algorithms dropped from SAML and OIDC IAM configs

**DEPRECATED**

65

New and updated SAML IAM configurations no longer support SHA-1; supported signature and digest algorithm values are now **SHA-256** and **SHA-512**. The unsigned JWT algorithm ( **none** ) is no longer available as a supported ID token signing algorithm in OpenID Connect configurations. Existing SAML configurations using SHA-1 continue to work for now but must migrate before a future enforcement step.

– Names exact algorithm values and migration note, no config key path. [snapshot-20261016](#)

### ←→ **BREAKING ON UPGRADE**

- ! Tag-driven domain update API endpoints now return **202 Accepted** instead of **200 OK**; callers that assume the update is complete when the request returns will behave incorrectly and must be updated to handle the async response.
- ! New or updated SAML IAM configurations can no longer use SHA-1 for signature or digest algorithms; existing configurations using SHA-1 (explicitly or via the previous SHA-1 default) continue to work until further notice but must migrate before a future enforcement step.

WAS THIS USEFUL?



# Clevis

[Feedback](#)[v24](#)[CODE ↗](#)[RANK](#)

Automated Encryption Framework

Clevis v24 adds mutual TLS authentication to the tang pin and ships several dracut boot-time fixes, including removing TPM1's network dependency and improving PKCS#11 and tang unlock reliability during early boot.

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

## 01 **Dracut boot-time module improvements** IMPROVED 68

Several fixes to Clevis's dracut modules: the TPM1 dracut module no longer requires network availability at boot (network dependency dropped); the dracut initrd PKCS#11 module now includes `libpcsclite_real.so` to enable PKCS#11 token access during early boot on distributions that ship the library; and `clevis-luks-askpass` is now ordered after `network-online.target` in the tang dracut module when systemd is included, ensuring network readiness before a tang unlock attempt.

– Names exact files and systemd ordering target for each fix v24

## 02 **Force-password bypass for LUKS key changes** IMPROVED 65

Clevis now passes `--force-password` to `luksChangeKey` to bypass pwquality/cracklib checks when setting Clevis-generated random keys.

– Exact flag and command named, mechanism explained v24

THINNER COVERAGE BELOW

## 03 **Optional mTLS support for tang pin** NEW 45

The tang pin now supports optional mutual TLS (mTLS), allowing authenticated connections to tang servers during unlock.

– Names the mechanism but no config keys or flags given v24

WAS THIS USEFUL?



# SPIFFE spire

Feedback ⓘ v1.15.4 NOTES RANK

The SPIFFE Runtime Environment

| SPIRE v1.15.4 adds configurable TLS policies for TLS listeners.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Configurable TLS policies for TLS listeners NEW 30

SPIRE now supports configuring TLS policies to be applied to TLS listeners.

- No config key, flag, or policy options named. v1.15.4

WAS THIS USEFUL?  

Aembit is a runtime IAM for AI agents and workloads. Enforce access policies, keep downstream credentials out of agents, audit access, and revoke it with a kill switch.

Aembit introduced MCP Tool Access Control, letting teams enforce per-tool allow/block policies on MCP servers with fail-closed defaults and full audit logging.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

## 01 MCP Tool Access Control in MCP Identity Gateway NEW

88

Introduces per-tool allow/block rules inside an access policy, controlling `tools/list` discovery visibility and `tools/call` invocation independently. Supports exact tool names and glob patterns (e.g., `get_user*`, `admin_*`, `list_*`, `delete_*`) so new tools inherit rules automatically, with rules enforceable in `Allow` mode (deny unmatched by default) or `Block` mode (permit unmatched by default), selectable per policy. Rules are scoped by blended identity — user identity, group membership, and workload/agent identity — so different principals sharing the same MCP server get different tool sets, and tool-filtering policy is evaluated before deeper inspection layers (e.g., CrowdStrike AIDR) in the enforcement chain. The gateway fails closed by default: a missing or malformed policy returns an empty `tools/list` and blocks `tools/call`, never failing open. A new `application.mcp.contentSecurity.toolFilter` audit event records which rule matched and why for every filtered `tools/list` response or blocked `tools/call` request.

Lock a read-only customer-support agent to only `list_*` tools on a GitHub MCP server, blocking all write and delete tools by glob pattern.

📌 In the Aembit console, open your access policy for the customer-support agent, enable MCP Tool Access Control, set mode to `Allow`, and add the rule: Tool pattern: `list_*` Leave all other tools unmatched — they are denied automatically. Save the policy; no MCP client or server changes are needed.

— Names config keys, modes, glob patterns and exact enforcement order.

Introducing MCP Tool Filtering: Fine-Gr...

WAS THIS USEFUL?



## ◆ Identity & Access Management

# Ping Identity

Feedback



15 Oct 2026 (2026-10-15)

BLOG ↗

RANK

Ping Identity is an identity and access management platform providing authentication, single sign-on, multi-factor authentication, and identity governance services.

Ping Identity's rapid channel update refines Advanced Reporting, making report parameters optional and switching entitlement review usernames from GUIDs to human-readable form — the latter a breaking change for dependent integrations.

↳ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

### 01 Entitlement review 'Completed By' now shows usernames

62

BREAKING

In Advanced Reporting, the **Completed By** field in entitlement review items now returns usernames instead of GUIDs. Saved reports or integrations that expect GUIDs must be updated to account for the change.

– Names exact field and gives clear before/after plus migration note.

15 Oct 2026 (2026-10-15)

THINNER COVERAGE BELOW

### 02 Optional parameters in Advanced Reporting

IMPROVED

48

In Advanced Reporting, parameters such as **status**, **journeys**, and **roles** are now optional in several out-of-the-box reports, allowing reports to be run without selecting values that aren't needed.

– Names the parameters but no exact report list or UI steps.

15 Oct 2026 (2026-10-15)

↳ BREAKING ON UPGRADE

! The **Completed By** field in entitlement review items now returns usernames instead of GUIDs — saved reports or integrations that expect GUIDs must be updated.

WAS THIS USEFUL?



# Clerk

Feedback



Clerk MCP server (2026-10-15)

BLOG ↗

RANK

Clerk is an authentication and user-management platform that provides pre-built UI components, APIs, and SDKs for adding sign-up, sign-in, and identity features to applications.

Clerk launched a revamped MCP server that lets AI agents manage Clerk applications, users, configuration, and production setup directly.

WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

## 01 Revamped Clerk MCP server NEW 37

A new Clerk MCP server enables managing Clerk applications, users, configuration, and production setup directly from an AI agent.

– Names capability but no setup steps, commands, or scope limits

Clerk MCP server (2026-10-15)

WAS THIS USEFUL?



Discover the latest features, improvements, and updates to Ory.

Ory refined latency observability in Hydra with finer-grained histogram buckets for its request duration metrics.

WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Finer-grained latency histogram buckets in Hydra IMPROVED

50

Ory Hydra now reports finer-grained latency histogram buckets for the `http_requests_duration_seconds` and `http_response_time_seconds` metrics, improving the resolution of request latency measurements.

- Names exact metrics but no concrete bucket values or usage steps.

Ory Hydra, Ory Kratos, Ory Keto, Ory Oa...

WAS THIS USEFUL?



# DigiCert CertCentral

Feedback



October 20, 2026

NOTES

RANK

DigiCert CertCentral is a web-based platform for ordering, issuing, managing, and renewing SSL/TLS certificates and other digital trust credentials.

DigiCert CertCentral clarified how its default intermediate chain is surfaced in product settings and removed unneeded self-signed roots from Apache certificate downloads for G5 and G2 TLS certificates.

← WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

## 01 Apache download drops self-signed G5/G2 roots

BREAKING

71

Apache `.crt` download files for G5 and G2 TLS certificates now contain only the intermediate CA and cross-signed root — the self-signed G5 and G2 roots are removed from `DigiCertCA.crt`. As of October 20, 2026 this drops the file from three certificates to two; scripts that remove a certificate by position or expect a specific certificate count will break.

— Names exact file and before/after count, but no migration command given

October 20, 2026

## 02 DigiCert default option in intermediate chain menu

IMPROVED

62

The `Default intermediate chain` menu in `Settings > Product Settings` now shows a `DigiCert default (recommended)` option so admins can distinguish DigiCert's rolling default chain from a chain they explicitly selected and pinned.

Verify your product is pinned to a specific intermediate chain — not the rolling DigiCert default — before the November 12, 2026 G5 migration changes it.

📌 In CertCentral, go to Settings > Product Settings. Select the certificate product. In the `Default intermediate chain` menu, check whether it shows `DigiCert default (recommended)`. If it does and you need a fixed chain, select your preferred chain and click `Save Settings`.

— Exact UI path and option name given with verification steps

October 20, 2026

← ALSO FROM THESE RELEASES

•  
•  
•

## ↳ BREAKING ON UPGRADE

! As of October 20, 2026, the Apache `DigiCertCA.crt` file for G5 and G2 TLS certificates drops from three certificates to two (intermediate CA and cross-signed root only). Scripts that remove a certificate by position or that expect a specific certificate count will break.

WAS THIS USEFUL?



## GOVERNANCE, RISK & COMPLIANCE

### ◆ Integrated & IT Risk Management

## LogicGate

Feedback



Introducing LogicGate's Cyber Risk Management Agents (2026-10-15)

**BLOG** ,

LogicGate Risk Cloud is the leading AI-powered GRC platform with 30+ applications for enterprise risk, third-party risk, cyber risk and compliance management.

LogicGate launches a Cyber Risk Management Agentic Application with three AI agents — Triage, Assessment, and Insights — for automated first-pass risk workflows.

- › Adds the Cyber Risk Triage Agent, which evaluates incoming risk records against your governance definitions, separates true cyber risks from routine issues or policy exceptions, and routes invalid submissions back to the original submitter with guidance.
- › Adds the Cyber Risk Assessment Agent, which synthesizes linked threat data, evaluates mitigating IT controls, and proposes justified risk scores with auto-filled remediation plans aligned to your organization's framework.
- › Adds the Insights Agent, which transforms complex risk data into a plain-language chat interface, generates executive-ready reports, and dynamically flags escalating technical threats requiring immediate attention.

### COVERAGE

**NEW** compliance framework

[NIST Cybersecurity Framework and ISO/IEC 27001 control mapping in Cyber Risk Assessment Agent](#)

[All of LogicGate's coverage changes →](#)

WAS THIS USEFUL?



---

**AGENT**

# Runlayer

Feedback



1.28.1407 - Detect agents inside IDEs (2026-10-16)

BLOG ↗

RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

Runlayer's Detect module now surfaces IDE-native AI agents and coding-assistant extensions that previously evaded inventory, and connector policies gained finer per-tool control.

↳ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

## 01 IDE-native agent and extension detection

NEW

75

Runlayer Detect now inventories IDE-native AI agents that were previously invisible because they don't read the standalone Claude Code, Codex, or VS Code config files — including Claude and Codex agents in Xcode 26.3+ (read from its own `CodingAssistant` folder), Visual Studio Copilot agent mode, JetBrains AI Assistant, Gemini in Android Studio, GitHub Copilot for Xcode and Eclipse, Firebender, Positron, and PearAI. It also recognizes Codex for VS Code, CodeGPT, Qoder CN, Bito, Zencoder, Tabby, twinnny, Verdent, Pieces, and the Firebender JetBrains plugin as AI extensions in the Detect inventory.

Verify that Xcode 26.3+ `CodingAssistant` MCP servers are now visible in your Detect inventory after upgrading — agents in that folder were previously missed entirely.

📌 In the Runlayer console, go to Detect and filter by IDE to confirm Claude and Codex agents running from the Xcode `CodingAssistant` folder appear in the agent inventory.

– Names every agent/IDE and explains why detection was missed before

1.28.1407 - Detect agents inside IDEs (...)

THINNER COVERAGE BELOW

## 02 Per-tool granularity for connector policies

IMPROVED

15

Connector policies gain per-tool granularity, per the release summary; no further mechanism or configuration detail is given.

– Mentioned only in summary with no mechanism or scope

1.28.1407 - Detect agents inside IDEs (...)

## COVERAGE

**NEW** detections

[AI Watch Detect: 22 open-source coding harnesses](#), [AI Watch Detect: second wave of vendor coding agents with China-based default endpoints](#)

**UPDATED** integration [Looker and LaunchDarkly MCP connectors — generally available](#)

[All of Runlayer's coverage changes →](#)

### 1.28.1396 - Gateway spend tracking and session stop (2026-10-15)

**BLOG** ↗

**STABLE**

Adds session stop API, [LLM Gateway](#) spend tab, connector approval intake, idle-connector keep-alive, upstream User-Agent identification, and expanded Runlayer Plugin artifact operations.

- › Adds a **Session stop** API for Security, IT, and Super Admins to halt a live agent session with a reason; subsequent [MCP](#) tool calls carrying that session's **id** are denied until an admin clears the stop, with each stop, first denied call, and clear recorded in the audit log.
- › Adds a **Settings → LLM Gateway → Spend** tab listing every LLM call the gateway reported, with cost, tokens, and charged usage rule, and totals spend by agent run, device, or model over a chosen time range; calls the gateway cannot price show as **unpriced** rather than **\$0**.
- › Adds connector approval intake so workspace admins can direct employees to an external approval form with custom instructions, with the option to disable native connector submissions.
- › Runlayer Plugin agents can now publish, find, read, edit, and delete artifacts, and delete skills and agents directly, without a separate Runlayer MCP server connection.
- › Runlayer now names the AI client in the **User-Agent** sent to Streaming HTTP and SSE connectors and identifies itself as Runlayer on OAuth requests, enabling out-of-the-box connectivity to servers that check the calling client or reject OAuth requests without a **User-Agent**.

[+1 more](#)

### 1.28.1381 - Secret stores and agent eval scoring (2026-10-15)

**BLOG** ↗

**STABLE**

Adds secret stores (AWS Secrets Manager), agent eval scoring, Terraform-managed scanner/AI Watch defaults, plugin-scoped tokens, and hardened connector [sandboxing](#).

- › Adds secret stores: keep connector credentials in your own AWS Secrets Manager and reference them in connector headers and environment variables; workspace admins manage stores under Settings.
- › Adds `Matches expected output` eval scorer for agent runs — save the expected answer on a reviewed run, then replay it to regression-test prompt and model changes.
- › Adds `runlayer_security_scanner_settings` Terraform resource to manage organization-wide security scanner defaults as code; self-hosted deployments require backend 1.28.1220 or later.
- › Adds `runlayer_ai_watch_enforce_entry` Terraform resource to manage each AI Watch Enforce allowlist or Protect denylist entry (remote `MCP` URL, stdio package, or built-in) as code; self-hosted deployments require backend 1.28.1171 or later.
- › Adds `PATCH /ai-watch/policy-rules/{id}` to edit the organization-wide AI Watch policy, now an editable default ( `Monitor` with `Sessions` off) rather than a computed value; the rule list shows which deployments differ from it.

+7 more

#### ↳ BREAKING ON UPGRADE

- ! The gateway refuses to start or dial connectors configured for local deployment on any proxy path; a direct `POST /api/v1/proxy/{id}/mcp` call to one is recorded as a security violation and can end the caller's session.
- ! `stdio` connectors that carry `LD_*`, `DYLD_*`, or `GCONV_PATH` variables in their environment or inherit list will be refused at launch.
- ! Hosted `stdio` connectors ( `npx`, `uvx`, `uv` ) that pass file-reading options ( `--env-file`, `--userconfig` ), filesystem path values, or disallowed environment variable names ( `HOME`, `PATH`, `PREFIX`, `XDG_*`, `NODE_PATH`, `PYTHONPATH` ) are refused on save and at launch.

WAS THIS USEFUL?



## DEPLOY

# Portkey

[Feedback](#)[changelog · Oct 16, 2026](#)[CHANGELOG ↗](#)[RANK](#)

Democratize and productionize Gen AI across your entire org with Portkey

Portkey's latest changelog adds a new Decisions API for typed-judgment models, a startHooks lifecycle stage for pre-auth traffic screening, MCP Gateway guardrails, and expands provider coverage with Meta, ElevenLabs, and xAI image generation.

↳ **WHAT SHIPPED · 10 FEATURES** most completely described first [? what's the number?](#)

## 01 Decisions API for typed judgments NEW 95

A new Decisions API endpoint lets callers send a question with structured state and receive a typed judgment (boolean, choice, or score) with probabilities and a calibrated confidence score. Requests route through the full gateway stack (auth, routing, retries, caching, logging, cost tracking); initial support covers TypeSafe's **Jev** model family.

Call the Decisions API to get a calibrated boolean judgment from a TypeSafe Jev model — useful for automated triage or policy enforcement pipelines.

```
$ curl -X POST https://api.portkey.ai/v1/decisions \
-H 'x-portkey-api-key: <your-api-key>' \
-H 'Content-Type: application/json' \
-d '{
  "model": "jev",
  "question": "Does this request contain PII?",
  "state": {"text": "My name is Jane Doe and my SSN is 123-45-6789"},
  "output_type": "boolean"
}'
```

– Runnable curl example against named endpoint with model and output types.

[changelog-20261016-bffacb2a](#)

## 02 startHooks lifecycle stage with Header Check/Transform NEW 85

A new **startHooks** lifecycle stage screens and normalizes traffic before authentication, routing, or other hooks run. It introduces **Header Check** (allow/block requests by header key-value pairs) and **Header Transform** (remove, set, or rewrite headers before routing, with **x-portkey-\*** headers protected from modification); the

existing `allowedRequestTypes`, webhook, and `JWT` header lookup checks can now also run as `startHooks`.

– Names every check and header pattern but no config snippet. `changelog-20261016-bffacb2a`

### 03 New provider support: Meta and ElevenLabs NEW 70

Adds `Meta` as a provider with support for `Chat Completions`, `Messages API`, `Responses API`, and `Files API`, and adds `ElevenLabs` as a provider with support for text-to-speech and speech-to-text, both with full observability and reliability features.

– Names both providers and their supported APIs, no setup steps.

`changelog-20261016-bffacb2a`

### 04 Anthropic-protocol auto-detection on `/v1/models` NEW 65

`/v1/models` now auto-detects the Anthropic protocol so Claude Code and other Anthropic-native clients receive Anthropic's native model shape instead of the OpenAI-normalized shape.

– Named endpoint, no sample request shown. `changelog-20261016-bffacb2a`

### 05 MCP Gateway guardrails for tool calls IMPROVED 60

MCP Gateway now runs MCP tool calls through the same pre- and post-execution guardrail checks already available for `LLM` requests; denied calls return a structured error that includes the guardrail and check ID.

– Mechanism described but no endpoint or config to act on. `changelog-20261016-bffacb2a`

THINNER COVERAGE BELOW

### 06 xAI image generation support IMPROVED 55

The xAI provider now supports image generation using OpenAI-compatible parameters plus xAI-specific aspect ratio and resolution controls.

– Names the controls but no example call. `changelog-20261016-bffacb2a`

### 07 Multi-JWKS JWT authentication IMPROVED 50

JWT authentication now accepts multiple `JWKS` URLs and attributes incoming requests to the matching workspace user.

– Mechanism named but no config key given. `changelog-20261016-bffacb2a`

### 08 Per-attempt retry logging IMPROVED 45

Each retry attempt now gets its own log row with its own status and response, instead of the log surfacing only the final result after retries complete.

– Clear before/after but no named surface to use directly. `changelog-20261016-bffacb2a`

## 09 Conditional routing on multipart form data IMPROVED

45

Conditional routing can now read fields from `multipart/form-data` requests, such as image edits and audio transcriptions.

– Use cases named, no field or flag specified. `changelog-20261016-bffacb2a`

## 10 Streaming thinking blocks in Messages API NEW

45

Reasoning models accessed through the Messages API now stream their thinking phase as a `thinking` content block.

– Names the content block but no usage example. `changelog-20261016-bffacb2a`

WAS THIS USEFUL?



---

GOVERN

# Zenity

[Feedback](#)[Zenity Launches Universal Runtime Control for Claude Enterprise](#)[NOTES](#)[RANK](#)

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

Zenity shipped universal runtime control for Claude Enterprise, enforcing pre-execution allow/deny decisions on tool calls across every Claude surface via Anthropic's Inference Hooks, with no endpoint agent required, alongside new Boundaries capabilities for read/write classification, tool provenance, prompt-injection session blocking, and denial visibility.

→ WHAT SHIPPED • 5 FEATURES most completely described first [? what's the number?](#)

## 01 Read/write classification in Zenity Boundaries NEW

88

Zenity Boundaries can classify supported tool calls as read or write, letting policy allow reads while blocking deletes, bulk updates, external shares, shell commands, or form submissions that violate policy, without depending on tool names that vary across connectors. This is enabled via **Validate tool calls** on a Boundary scoped to Claude Enterprise, classifying the target connector as read-only and setting the write/delete action to Prevent.

Enable **Validate tool calls** with a Boundary that allows reads but blocks write/delete operations on shared drives — stopping deletions before they execute with no agent deployed to endpoints.

📌 In the Zenity console, go to Zenity Boundaries > New Boundary, set scope to Claude Enterprise, enable **Validate tool calls**, classify the target connector as read-only, and set the write/delete action to Prevent.

– Exact console steps given for configuring the boundary.

[Zenity Launches Universal Runtime Contr...](#)

## 02 Tool provenance enforcement for MCP servers NEW

88

Boundaries can restrict tool calls to verified MCP servers and connectors, or block calls from unverified or local origins, using provider metadata supplied with each **tool\_call** event. Configured under Zenity Boundaries > Tool Provenance by setting **Third-party MCP servers** to **Verified only** and unverified/local origins to Prevent, applying organization-wide to Claude Code without an endpoint hook.

Restrict Claude Code's MCP server access to only Anthropic-verified servers organization-wide, blocking calls from unverified or local origins before execution.

 In the Zenity console, go to Zenity Boundaries > Tool Provenance, set **Third-party MCP servers** to **Verified only**, and set unverified/local origins to Prevent — applies across Claude Code without any endpoint hook.

– Exact console navigation and setting values named.

Zenity Launches Universal Runtime Contr...

### 03 Runtime enforcement via Anthropic Inference Hooks NEW

75

Zenity enforces pre-execution allow-or-deny verdicts on Claude tool calls using Anthropic's **tool\_call** Inference Hooks event, blocking prohibited operations before they run across claude.ai (web, desktop, mobile), Cowork, Claude Code, Claude for Microsoft 365, Claude in Chrome, Claude's built-in browser, and Claude Tag in Slack — with no endpoint agent or browser extension required.

– Mechanism and scope named but no setup steps for this specific capability.

Zenity Launches Universal Runtime Contr...

### 04 Denial visibility in Activity Feed and audit log NEW

75

Blocked tool-call denials appear in Claude's Activity Feed with a reference to the associated finding, giving security teams a reviewable chain of events per blocked action; reviewers can filter by **Denied** in the Activity Feed and follow the Compliance API audit link to see the injected instruction, tainted session marker, and tool call parameters.

Review a blocked prompt-injection-influenced tool call and its full event chain after a Claude Tag denial in a shared Slack channel.

 In Zenity, open Activity Feed > filter by **Denied**, select the blocked Claude Tag event, and follow the Compliance API audit link to see the injected instruction, the tainted session marker, and the associated tool call parameters.

– Exact filter and navigation steps given for reviewing denials.

Zenity Launches Universal Runtime Contr...

### 05 Prompt-injection tainted session blocking NEW

65

When Zenity detects an injected instruction in content Claude reads, it marks the session as tainted and blocks any subsequent tool call that acts on that instruction, surfacing the denial reason in Claude with the full event chain available in Zenity's Compliance API audit log.

– Mechanism explained but no concrete endpoint or config shown.

WAS THIS USEFUL?



## PLATFORMS & UTILITIES

### ◆ General Utilities & Libraries

## PyPA python-pipx

Feedback



1.18.2

NOTES ↗

RANK

The python-pipx CLI installs and runs Python applications in isolated environments while exposing their commands on the user's PATH.

pipx 1.18.2 adds support for installing local packages via absolute file URLs.

←→ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

### 01 Absolute file URL support for local installs NEW 70

`pipx install` now supports absolute file URLs (e.g. `file:///path/to/package`) for installing local packages, in addition to bare paths.

Install a locally built wheel or source tree using an absolute file URL instead of a bare path.

```
$ pipx install file:///home/user/projects/mytool
```

– Names exact flag/command and URL format with runnable example 1.18.2

WAS THIS USEFUL?



// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 16, 2026

Releases &amp; features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe

Feedback on this issue