

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 20, 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 208 changes](#)Do you prefer **scroll**?[▶ // CONTENTS](#) [SHOW](#)

54 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. Aembit can now see live MCP sessions and kill them mid-flight, closing the gap where an agent's vim ? help outlives its mandate. Kasten K10 publishes SBOM and OpenVEX

with its images, so scanner noise can be suppressed with evidence rather than a blanket exception. StepSecurity inventories Claude Code and Codex plugins published across an org's repos.

RESPOND

Cut off a running MCP agent session instead of waiting for the token to expire

Agent-to-tool traffic over MCP has been largely a black box: once a workload holds a credential, there is no session to look at and nothing to terminate short of rotating the secret. Session visibility plus revocation gives an on-call responder a containment action for agent activity, and schema-validated tool outputs reject malformed responses before a model acts on them.

[Aembit](#)

IDENTIFY

Find the AI that nobody registered — local models on endpoints, agent plugins in repos

Shadow AI has moved past SaaS chatbots into models running on laptops and agent plugins committed to org repos, neither of which show up in an IdP or proxy log. Backslash discovers local models on endpoints, StepSecurity enumerates and audits published Claude Code and Codex plugins, and 1Password's SaaS Manager now tracks GitHub Copilot usage and exports AI consumption to CSV — the raw material for an AI-use policy that references real inventory.

[Backslash](#) · [StepSecurity](#) · [1Password](#)

PROTECT

Suppress a CVE against the vendor's own evidence, not a blanket scanner exception

Kasten K10 9.0.8 now publishes SBOM and OpenVEX alongside its images, so a scanner hit on a bundled dependency can be marked not-affected with the upstream's attestation rather than muted by hand every release. Chainguard's Custom Assembly Overlays let a team pin per-tag image customizations instead of maintaining a fork of a hardened base.

[Kasten K10](#) · [Chainguard](#)

Bugcrowd

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

RANK

Bugcrowd is a crowdsourced security platform that connects organizations with ethical hackers to find and report vulnerabilities through bug bounty and penetration testing programs.

Bugcrowd added API support for programmatically triggering Savant Pathseeker Tests.

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

01 API-triggered Savant Pathseeker Tests NEW

45

Bugcrowd now enables triggering Savant Pathseeker Tests via the API, allowing automated or pipeline-driven test execution.

– Names the API capability but no endpoint or parameters given

changelog-20261020-3f5ccbd4

WAS THIS USEFUL?



◆ **Reconnaissance & OSINT**

nerva

Feedback



v1.72.0

NOTES ↗

Fast service fingerprinting CLI for 170+ protocols (TCP/UDP/SCTP) - built by Praetorian

↳ **GET THIS VERSION**

```
$ git clone --branch v1.72.0 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.72.0
```

COVERAGE

NEW

detection [Azure App Service Kudu SCM fingerprinter](#)

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

WAS THIS USEFUL?



APPLICATION SECURITY

Levo

Feedback



changelog · Oct 20, 2026

CHANGELOG

RANK

Levo continuous API Security Assurance  observe, detect, and protect your APIs.

Levo's 1.3.x release adds AI-assisted triage for bulk vulnerability selections, a gateway path-bypass authentication test, and a set of AI agent discovery and observability features, alongside multi-GitHub-account linking, Streamable HTTP support for its MCP server, AI-driven login for DAST scans, and deeper source-code API discovery.

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

01 AI agent discovery and observability features NEW

60

Levo now discovers AI agents, models, and MCP tools from cloud providers, adds a next-generation AI traffic tagger that attributes traffic to individual agents and sessions, and provides a full Agent Detail view for inspecting discovered AI agents.

- Three related AI agent additions described but no config or API detail

changelog-20261020-c5ed4996

THINNER COVERAGE BELOW

02 AI-guided login for DAST scans NEW

45

Adds AI-driven authentication (AI-guided login with sign-in verification) for DAST web application scanning.

- Explains mechanism but no config or scan setup detail

changelog-20261020-c5ed4996

03 Authentication-bypass test for gateway path handling NEW

40

Adds an authentication-bypass test targeting gateway path-handling weaknesses.

- Names the test type but not how to configure or run it

changelog-20261020-c5ed4996

04 Streamable HTTP transport for MCP server NEW

40

Introduces Streamable HTTP transport support for the MCP server.

- Names the transport and server but gives no usage detail

changelog-20261020-c5ed4996

05 AI-assisted triage for bulk vulnerability selections NEW

35

Adds AI-assisted triage to bulk vulnerability selections, letting teams assess many findings at once.

– States the capability but not how triage is invoked or scored

changelog-20261020-c5ed4996

06 Multi-GitHub-account linking NEW 30

Adds support for linking multiple GitHub accounts.

– Only named in summary, no further detail on setup

changelog-20261020-c5ed4996

07 Deeper API discovery from source code IMPROVED 15

Deepens API discovery from source code.

– Single vague line with no specifics on scope or mechanism

changelog-20261020-c5ed4996

WAS THIS USEFUL?



Fastly

Feedback



Retrieve historical metrics with improved performance and flexibility (2026-10-20)

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 introduced Platform Metrics V2, a historical metrics API with cursor-based pagination, predefined metric sets, and flexible grouping options.

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

01 Platform Metrics V2 historical metrics API NEW 65

Fastly added new Platform Metrics API endpoints for retrieving historical metrics data, using cursor-based pagination to handle large datasets. The API supports predefined metric sets that bundle related metrics for simpler querying, and lets users group historical metrics by point of presence or datacenter.

– Describes mechanism and options but no exact endpoint paths given.

Retrieve historical metrics with improv...

WAS THIS USEFUL?



GitGuardian SaaS

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

RANK

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

GitGuardian SaaS added in-network validity and permissions checks for vaulted secrets via ggscout, surfaced the results across the NHI inventory, introduced a new check-run skip API endpoint, reprioritized on-demand Jira Cloud scans, and renamed its subscription plans.

↳ WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Validity and permissions checks for vaulted secrets NEW

93

ggscout 0.34.1 adds in-network validity checks and permissions analysis for vaulted and Kubernetes secrets, running after classification and before hashing so internal vaults, VPN-gated databases, and provider-restricted endpoints are reachable without allowlisting GitGuardian IPs; deployments can opt out with a new `[check] enabled = false` config key. Results surface in the NHI inventory as a **Validity** column and filter (**Valid** / **Invalid**), token permission scopes (e.g. `repo` , `admin:org` on GitHub; `api` on GitLab) per credential in the NHI identity detail, plus owner account, expiry date, and provider metadata added via **ggscout** ; validity verdicts from vaulted or Kubernetes secrets merge with any matching secret incident, keeping the most recent verdict across both surfaces.

Opt out of validity and permissions checks on a specific ggscout deployment where outbound provider calls are not permitted.

```
ini

[check]
enabled = false
```

– Config key, exact columns and scopes named with mechanism explained

changelog-20261020-b4f8e432

THINNER COVERAGE BELOW

02 On-demand scan prioritization for Jira and API launches

IMPROVED

50

On-demand Jira Cloud scans launched from the dashboard or the public API now queue ahead of recurrent and system scans; public API launches using a service account token

on AWS S3 and version control sources now also count as on-demand scans.

– Clear behavior change but no config or command to invoke it changelog-20261020-b4f8e432

03 **Valid saved view for discovered secrets** NEW 45

Adds a built-in **Valid** saved view to the Endpoint Protection discovered secrets page, listing all secrets currently checked as valid.

– Names exact UI location but no further mechanism changelog-20261020-b4f8e432

04 **Check-run skip API endpoint** NEW 45

Adds a new **POST /v1/check-runs/{check_run_id}/skip** API endpoint.

– Exact endpoint named but no explanation of behavior changelog-20261020-b4f8e432

05 **Renamed plans and plan comparison page** IMPROVED 40

GitGuardian plans are renamed to Free (Starter), Growth, and Enterprise, and the **Plan and usage** page now compares features per plan grouped as on the pricing page.

– Names the page and new plan names, no deeper mechanism changelog-20261020-b4f8e432

↳ ALSO FROM THESE RELEASES

:

WAS THIS USEFUL?



Truffle Security trufflehog

[Feedback](#)[v3.99.3](#)[NOTES](#)[CODE](#)

TruffleHog finds, verifies, and analyzes leaked credentials across Git, chat, wikis, cloud storage, and more

Adds per-logger debug level override for source logging, giving operators fine-grained verbosity control per source.

[GET THIS VERSION](#)

```
$ git clone --branch v3.99.3
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.99.3
```

- › Adds per-logger debug level override for source logging, allowing fine-grained debug verbosity control per source.

COVERAGE

[NEW](#) detection [SonarQube Server token detector](#)

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

WAS THIS USEFUL?



Veracode

Feedback



changelog · Oct 20, 2026

CHANGELOG

RANK

Veracode is a cloud-based application security platform that scans code, binaries, and running applications to find and help remediate software vulnerabilities.

Veracode's CLI gains native static-scan commands and C# source packaging, GitHub Workflow Integration adds runner configuration for SCA scans, and EASM gets a redesigned interface.

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

01 New scan commands and C# source packaging in CLI NEW

85

Veracode CLI v2.53.0 adds `veracode scan static` to run static scans and review security findings within a Git workflow, `veracode findings static` to access cached static scan findings in the local Veracode directory, a `--csharp-source` flag to package C# source code directly for Static Analysis without compiling it into binaries, and now automatically includes `.properties` files when packaging Java source code.

– Names exact CLI commands and flags, no usage walkthrough given.

changelog-20261020-369d3ca8

THINNER COVERAGE BELOW

02 build_runs_on support for SCA in GitHub Workflow Integration

NEW

50

GitHub Workflow Integration v2.1.9 adds `build_runs_on` support for SCA scans.

– Names the config key but not its values or effect.

changelog-20261020-369d3ca8

03 Redesigned EASM interface IMPROVED

20

EASM gains a redesigned interface with improved navigation and usability.

– Generic description with no specifics on what changed.

changelog-20261020-369d3ca8

WAS THIS USEFUL?



Semgrep

[Feedback](#)[snapshot-20261020](#)[NOTES](#)[RANK](#)

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

Semgrep's findings API now exposes Autofix PR and failure data for every deployment, while Multimodal gains an Automatic AI provider mode, Guardian adds Codex support, Agentic Workflows scheduling becomes admin-only, and GitHub App webhook secret rotation now propagates to all existing webhooks.

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

01 Autofix PR and failure data in findings API **IMPROVED**

90

The `GET /api/v1/deployments/{deploymentSlug}/findings` endpoint now returns `click_to_fix_prs` (pull requests created by Autofix) and `click_to_fix_failures` (failed Autofix attempts) on all Semgrep Code and AI-powered detection findings for every deployment, not only Autofix-beta deployments.

Pull Autofix PR and failure data for all findings in a deployment to audit which issues have been auto-remediated and which attempts failed.

```
$ curl -H 'Authorization: Bearer <token>'
'https://semgrep.dev/api/v1/deployments/<deploymentSlug>/findings'
| jq '[.findings[] | {id: .id, fix_prs: .click_to_fix_prs,
fix_failures: .click_to_fix_failures}]'
```

– Named endpoint, fields, scope change, and runnable example [snapshot-20261020](#)

02 Automatic AI provider mode for Multimodal **NEW**

65

Semgrep Multimodal AI provider settings now offer Automatic mode, which lets Semgrep manage and update the provider list (currently Amazon Bedrock, Cloudflare Workers AI, and OpenAI) without manual configuration; Automatic mode is the default when Semgrep Multimodal is first enabled.

– Names providers and default but no exact config path [snapshot-20261020](#)

03 Codex support in Semgrep Guardian **NEW**

65

Semgrep Guardian now supports Codex, enabling it to scan with Semgrep Code and Supply Chain via a Semgrep-hosted server (no Semgrep CLI required) using the fixed `guardian-default` ruleset.

– Names ruleset and mechanism but no setup steps [snapshot-20261020](#)

04 Admin-only Agentic Workflows schedule management

BREAKING

50

Agentic Workflows schedule and run management (create, change, delete, view schedules; trigger single runs, workflow packs, and bulk runs) is now restricted to organization admins via the API.

– States the restriction but not the affected endpoints [snapshot-20261020](#)

05 Webhook secret rotation propagation for GitHub App

IMPROVED

50

Rotating the webhook secret for a customer-owned GitHub App now also updates the secret on all existing webhooks, covering both GitHub Cloud and GitHub Enterprise.

– Describes behaviour change but no API or UI detail [snapshot-20261020](#)

WAS THIS USEFUL?



RANK

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

GitHub added a new REST API surface for managing Copilot model policies across enterprise, organization, team, and app scopes, with support for delegating control downward.

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

01 Copilot model policy API endpoints

NEW

92

New `GET` and `PATCH` endpoints let administrators retrieve and update Copilot model policies at multiple levels: `/enterprises/{enterprise}/copilot/model-policies` for enterprise-wide policy, `/enterprises/{enterprise}/teams/{team_slug}/copilot/model-policies` for per-enterprise-team policy, `/enterprises/{enterprise}/copilot/apps/model-policies` for enterprise apps, and `/orgs/{org}/copilot/model-policies` for organization-level policy. Enterprises can delegate model control downward to organizations via an `organization` mode, or to teams and apps via an `enterprise teams` mode, set by `PATCH` -ing the `mode` field.

Audit which Copilot models are currently permitted across your entire enterprise before rolling out a policy change.

```
$ curl -L \
  -H "Authorization: Bearer <token>" \
  -H "Accept: application/vnd.github+json" \
  https://api.github.com/enterprises/{enterprise}/copilot/model-policies
```

Delegate Copilot model control to individual organizations so each org can manage its own allowed models.

```
$ curl -L \  
-X PATCH \  
-H "Authorization: Bearer <token>" \  
-H "Accept: application/vnd.github+json" \  
-H "Content-Type: application/json" \  
https://api.github.com/enterprises/{enterprise}/copilot/model-  
policies \  
-d '{"mode": "organization"}'
```

– Names all endpoints, methods, and delegation modes with runnable examples.

Copilot model policy API for enterprise...

WAS THIS USEFUL?



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

GitLab's biggest move this window is the launch of the Duo Agent Platform — automated code review, CI/CD pipeline fixes, and natural-language Agentic Chat — alongside smaller Runner 19.5 observability fixes.

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

01 Fix CI/CD Pipeline Flow for failing pipelines NEW

88

Selecting `Fix pipeline with Duo` from a merge request's Overview or Pipelines tab has an agent read pipeline logs and propose inline code suggestions, open a new merge request with a fix, or post next-step comments; the same action is available via `Build > Pipelines` for pipelines not tied to a merge request, producing a new merge request with the proposed fix.

Send a failing pipeline to the agent for diagnosis and a proposed fix when the pipeline is not attached to a merge request.

 In the GitLab UI: Build > Pipelines > [select failing pipeline] > Fix pipeline with Duo

– Exact UI path and concrete agent outcomes described

Three developer tasks worth handing to ...

02 Code Review Flow for automated MR reviews NEW

85

Assigning `@GitLabDuo` as a reviewer or commenting `/assign_reviewer @GitLabDuo` on a merge request triggers automated diff, description, and related-file review comments posted inline; an automatic reviews option can also be enabled so every new merge request receives a first-pass review without manual assignment.

Trigger an automated first-pass code review on an open merge request without leaving the comment box.



```
$ /assign_reviewer @GitLabDuo
```

– Names exact slash command and reviewer behavior

Three developer tasks worth handing to ...

03 Agentic Chat, prompt library and session tracking NEW

83

Agentic Chat, opened via `GitLab Duo sidebar > Add new chat`, lets practitioners query merge requests, issues, and pipeline data in natural language; it's available in the GitLab UI, VS Code, JetBrains IDEs, and Visual Studio with the `Agentic` toggle enabled. An AI Resource Library ships ready-made prompts (e.g. `Summarize Recent Changes to MR/Issue`, `Debug Failing Pipeline`) for use with it, and an `AI > Sessions` view lets users follow agent reasoning in real time while it or Fix CI/CD Pipeline Flow runs.

– Names toggle, IDEs, prompt examples and sessions view

Three developer tasks worth handing to ...

04 Duo Agent Platform activation settings NEW

75

GitLab Duo Agent Platform is turned on for a group by enabling `Allow flow execution` and `Allow foundational flows` under `Settings > GitLab Duo > Change configuration > Flow execution`, which activates flows such as Code Review Flow and Fix CI/CD Pipeline Flow.

Activate the Duo Agent Platform flows for your group so Code Review Flow and Fix CI/CD Pipeline Flow are available to all members.

Settings > GitLab Duo > Change configuration > Flow execution: enable `Allow flow execution` and `Allow foundational flows` > select Code Review and Fix CI/CD Pipeline > Save changes

– Names exact settings path and toggle flags

Three developer tasks worth handing to ...

THINNER COVERAGE BELOW

05 CPU model logging and clearer Kubernetes OOM errors in Runner

IMPROVED

45

GitLab Runner 19.5 logs the `CPU` model of the job execution environment at job start for better observability of runner hardware, and surfaces a clearer Kubernetes executor error message when the helper container is `OOM-killed`.

– Describes two log/error improvements without config detail

product docs

WAS THIS USEFUL?



Mend.io Mend SAST

Feedback



Version 26.9.3 (19-October-2026)

NOTES ↗

Mend SAST is a static application security testing service that scans source code for vulnerabilities and helps teams prioritize remediation.

Gen 2 detection engine for Swift and PHP delivers improved CWE and framework coverage for more accurate, scalable scans.

- › Introduces the Gen 2 detection engine for Swift and PHP, with improved CWE and framework coverage for accurate, scalable scans.

COVERAGE

UPDATED

detection Gen 2 detection engine for Swift and PHP

All of Mend SAST's coverage changes →

WAS THIS USEFUL?



Mend.io Mend Platform

Feedback



Version 26.9.3 (18-October-2026)

NOTES

RANK

Mend platform is a traditional application security misses AI risk. Standalone AI security tools miss your code.

Mend Platform shipped aggregated multi-project SBOM export in CycloneDX and SPDX, took Custom Roles to general availability, and added an organization setting to exclude proprietary libraries from license-based workflow automation.

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

01 Aggregated SBOM export across projects NEW 72

Mend Platform now generates a single deduplicated SBOM across an entire application or multi-project scope, consolidating components into one CycloneDX or SPDX document. In the UI, navigate to the Application scope, select **Export SBOM**, choose **Aggregated** scope, pick CycloneDX or SPDX as the output format, and download the consolidated file.

Generate a single deduplicated SBOM across all projects in an application scope to feed into downstream tooling or compliance pipelines.

In the Mend Platform UI, navigate to your Application scope, select **Export SBOM**, choose **Aggregated** scope, select CycloneDX or SPDX format, and download the consolidated file.

Generate a single deduplicated SBOM across an entire application scope to feed into downstream compliance tooling.

In the Mend Platform UI, navigate to your Application, select **Export SBOM**, choose **Aggregated** scope, then select CycloneDX or SPDX as the output format and download the consolidated file.

— Names formats and exact UI navigation steps. Version 26.9.3 (18-October-2026)

02 Proprietary libraries excluded from license workflows NEW 70

A new organization setting, **Include Proprietary Libraries in Workflows**, excludes in-house proprietary components from **Empty License** automation workflow evaluations by default, preventing false-positive violations. Found in the UI under Organization Settings > Workflows, where the toggle can be disabled.

Prevent proprietary in-house libraries from triggering false-positive **Empty License** policy violations in automation workflows.

 In the Mend Platform UI, go to Organization Settings > Workflows, locate the **Include Proprietary Libraries in Workflows** setting, and disable it.

Prevent proprietary in-house components from triggering false-positive **Empty License** workflow violations across the organization.

 In the Mend Platform UI, go to Organization Settings, locate **Include Proprietary Libraries in Workflows** , and disable the toggle to exclude proprietary components from license-condition automation rules.

– Names exact setting and its UI location. **Version 26.9.3 (18-October-2026)**

THINNER COVERAGE BELOW

03 Custom Roles general availability **IMPROVED** 40

Custom Roles moved from closed beta to general availability, with an enhanced role breakdown view and a Create Role wizard for role management.

– Describes new UI elements but no navigation path given.

Version 26.9.3 (18-October-2026)

WAS THIS USEFUL?  

Mend sca is a traditional application security misses AI risk. Standalone AI security tools miss your code.

Mend SCA promoted Custom Roles to general availability with a streamlined role-creation flow, and added a settings toggle to keep proprietary libraries out of license-violation scans.

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

01 Proprietary library exclusion from license workflows NEW

80

In Organization Settings, disabling **Include Proprietary Libraries in Workflows** excludes in-house proprietary components from automated license-condition scans, preventing false-positive **Empty License** violations.

Prevent false-positive **Empty License** violations caused by proprietary libraries being swept into license-condition workflows.

In the Mend platform, go to Organization Settings and enable **Include Proprietary Libraries in Workflows** (set it to off/excluded) so in-house components are skipped during automated license scans.

- Exact setting and navigation path given with clear purpose product docs

THINNER COVERAGE BELOW

02 Custom Roles general availability NEW

45

Custom Roles moves from closed beta to general availability, with an enhanced role breakdown and a Create Role wizard for more intuitive role management.

- Names the wizard but no exact navigation or role details given product docs

COVERAGE

- NEW integration GitHub repository-level security advisories as an import source
- UPDATED supported platform Bun package manager GA in Mend CLI and repository integrations
- All of Mend SCA's coverage changes →

Adds aggregated CycloneDX/SPDX SBOM export for application and multi-project scopes, and an **Include Proprietary Libraries in Workflows** org setting.

- › Introduces an option to generate a single aggregated SBOM file for application and multi-project scopes, consolidating multiple project-level files into one document with deduplication and merged occurrence data; supports **CycloneDX** and **SPDX** standards.
- › Adds the **Include Proprietary Libraries in Workflows** organizational setting to control whether proprietary (in-house) components are included in automation workflows for **Empty License** conditions.

WAS THIS USEFUL?



Cloudsmith

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

RANK

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

Cloudsmith expanded upstream proxy support to the ActiveState Curated Catalog, added OIDC authentication for Buildkite and Google Cloud workloads, enabled listing packages from connected repositories, and extended cooldown policies to Ruby gems.

↳ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 **OIDC authentication for Buildkite and GCP** NEW 60

CLI v1.29.0 adds OIDC authentication for Buildkite pipelines and Google Cloud workloads, eliminating the need to store Cloudsmith API keys on build hosts.

– Names CLI version and benefit, no exact config/flag given. changelog-20261020-f257dc88

THINNER COVERAGE BELOW

02 **ActiveState Curated Catalog upstream proxying** NEW 55

Cloudsmith added ActiveState Curated Catalog upstream support, configurable via the Libraries tab, letting teams proxy and cache packages from the ActiveState catalog.

– Names config location (Libraries tab) but no further mechanism.

changelog-20261020-f257dc88

03 **Cooldown policies extended to Ruby gems** IMPROVED 50

Cooldown policies, which protect supply chains from newly published versions that may carry malware or lack community scrutiny, now extend to Ruby gem packages.

– Explains purpose and scope but no config detail. changelog-20261020-f257dc88

04 **List packages from connected repositories** NEW 40

CLI v1.29.0 adds the ability to list packages from connected repositories alongside your own packages.

– Brief description, no command or flag specified. changelog-20261020-f257dc88

WAS THIS USEFUL?



Socket

Feedback



Faster bulk ticket creation for manifest alerts (2026-10-20)

BLOG

RANK

Socket is a supply chain security platform that analyzes open-source package dependencies (npm, PyPI, etc.) to detect malicious code, risky behaviors, and vulnerabilities before developers install them.

Socket added bulk ticket creation for manifest alerts across multiple issue trackers and tightened OAuth access to enforce current role permissions on every request.

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

01 Bulk ticket creation for manifest alerts NEW 69

Users can bulk-create tickets in Jira, Linear, Asana, and ClickUp for manifest alerts (e.g. **Missing lockfile**), generating one ticket per repository and manifest file, scoped to the selected repository and branch. Shift-click range selection lets users select multiple alerts at once before creating tickets.

Select multiple manifest alerts across repositories and create one Jira ticket per repo/manifest in a single action.

In the Socket dashboard, navigate to your manifest alerts list. Shift-click to select the desired range of alerts, then choose **Create tickets** and select Jira (or Linear, Asana, or ClickUp) to generate one ticket per repository and manifest file.

Names exact trackers and selection mechanism with usage steps

Faster bulk ticket creation for manifes...

THINNER COVERAGE BELOW

02 OAuth access enforces current role permissions IMPROVED 53

For organizations with current-role checks enabled, OAuth token requests now validate both the token's granted scopes and the user's current role permissions, with role changes taking effect on the next request. Linked service-account OAuth requests now also check current service-account permissions alongside granted token scopes.

Explains mechanism and scope but no config surface or command

OAuth access stays in sync with role ch...

WAS THIS USEFUL?



Anchore syft

Feedback



v1.55.0

NOTES ↗

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

↳ GET THIS VERSION



```
$ git clone --branch v1.55.0 https://github.com/anchore/syft.git
# already have the repo? check out this version:
$ git checkout v1.55.0
```

COVERAGE

NEW

supported platform [IPK package scanning](#)

[All of syft's coverage changes →](#)

WAS THIS USEFUL?



StepSecurity is a security platform that hardens CI/CD pipelines, such as GitHub Actions, by detecting supply-chain risks, enforcing least-privilege runners, and monitoring workflow tampering.

StepSecurity introduces an Agent Plugins inventory that discovers and audits Claude Code and Codex plugins published across an organization's GitHub repositories.

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

01 Agent Plugins inventory for Claude Code and Codex NEW

85

A new **Agent Plugins** page under GitHub > AI Agents inventories Claude Code and Codex plugins detected from `.claude-plugin/plugin.json` or `.codex-plugin/plugin.json` manifests at the repository root or in subfolders. Each plugin shows declared version, description, author, license, homepage, and repository, plus the actual repository, branch, and commit StepSecurity scanned, with a direct link to the manifest. A **Used by** view lists which repositories in an org enable the plugin in their settings (making it available to all contributors there), the page lists every skill and other component a plugin supplies with hooks shown by the event they run on, and results can be filtered by agent or repository and searched by plugin or component name.

⋮

Audit which Claude Code or Codex plugins your org's repos publish and what hooks they register — navigate to GitHub > AI Agents > Agent Plugins in the StepSecurity console.

📌 In the StepSecurity console, go to GitHub > AI Agents > Agent Plugins. Use the **Filter by agent** dropdown to select **Claude Code** or **Codex**, then search by component name to find plugins that register hooks on a specific event.

Check which repos expose a specific plugin to all their contributors via the **Used by** view, to assess blast radius before modifying or removing a plugin.

📌 In the StepSecurity console, go to GitHub > AI Agents > Agent Plugins. Click a plugin name to open its details panel, then select the **Used by** tab to see every repository that has enabled it in its settings.

— Names manifests, fields, UI path and filters; no limits or numbers given

Agent Plugins Inventory for GitHub Repo...

WAS THIS USEFUL?



REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Finite State Platform

Feedback



October 19, 2026 - Finding history and license details (2026-10-19)

[BLOG](#)

RANK

Finite State Platform is a firmware and software supply chain security tool that analyzes binaries and embedded devices to identify vulnerabilities, SBOM components, and risks.

Finite State Platform added visibility into Finding history and declared license details for viewers, alongside expanded binary component detection across several embedded libraries.

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

01 Component detection for embedded libraries

NEW

50

Binary firmware scans now detect the `nanoMODBUS`, `nanoprintf`, `cifra`, `TinyCrypt`, and `MPLAB Harmony` libraries.

– Names specific libraries but no scan mechanism or usage path

October 16, 2026 - Component detection,...

02 Finding history visible to viewers

NEW

40

Finding history — including status updates and comments — is now exposed to any user who has view access to the Project or Folder, along with complete declared license details.

– Clear description but no UI path or mechanism detail

October 19, 2026 - Finding history and ...

WAS THIS USEFUL?



NETWORK SECURITY

Sophos Access Points

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

RANK

Sophos Access Points added support for configuring 6 GHz band steering through Sophos Fusion.

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

01 6 GHz band steering via Sophos Fusion NEW 40

Sophos Access Points now support configuring 6 GHz band steering, manageable from Sophos Fusion.

– Names feature and management console but no steps or mechanism

changelog-20261020-b31053b3

WAS THIS USEFUL?



Sublime Security

[Feedback](#)

Platform v1.90.8 (Flexible Threat Log Views) (2026-10-19)

[BLOG](#)[RANK](#)

Sublime Security is an email security platform that uses detection-as-code rules to identify and block phishing, business email compromise, and other malicious messages.

Sublime Security's Threat Log views gained more flexible filtering, letting analysts switch State and Verdict filters without leaving a view and spot unreviewed messages faster via a new count badge.

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

01 Adjustable filters and counts in Threat Log views IMPROVED

70

The **Needs Remediation** view's **State** filter is now adjustable, letting analysts switch to states like auto-remediated or mixed without leaving the view, while the **Needs Review** view's **Verdict** filter is now adjustable to surface messages with any verdict. **Needs Review** also gained a message count badge so unreviewed messages are always visible at a glance. Default filters were updated: **Needs Remediation** now opens on unreviewed Malicious messages with no remediation actions applied, and **Needs Review** opens on Spam, Unknown, and Suspicious messages.

Review all auto-remediated messages without leaving the **Needs Remediation** view by switching the State filter.

📌 In the Dashboard, go to Threat Log > Needs Remediation, then open the State filter and select **auto-remediated**.

Inspect Malicious-verdict messages directly from the **Needs Review** view by changing the Verdict filter.

📌 In the Dashboard, go to Threat Log > Needs Review, then open the Verdict filter and select **Malicious**.

— Names exact filters, views, states, and defaults with clear navigation steps.

Platform v1.90.8 (Flexible Threat Log V...

COVERAGE

NEW detectionsAttack Score Flash: Potentially Malicious (beta), Attack Score Flash: Spam (beta)

UPDATED detection

Attack Score Flash: Advanced Graymail Detection (renamed from Fuzzy Attack Score: Advanced Graymail Detection)

All of Sublime Security's coverage changes →

Platform v1.90.79 (Attack Score Flash - Public Beta) (2026-10-19)

BLOG ↗

STABLE

Attack Score Flash (public beta) adds `beta.attack_score_flash()` to all inbound mail, integrating verdicts into the existing Attack Score and ASA pipeline.

- › Adds `beta.attack_score_flash()` MQL function (Enterprise-only for malicious/potentially_malicious verdict; all tiers for spam) to evaluate all inbound mail and return verdicts of `malicious`, `potentially_malicious`, `spam`, or `graymail` — enabling detections that can catch messages Attack Score alone may miss.
- › Integrates Attack Score Flash verdicts into the existing Attack Score pipeline — `potentially_malicious` flags produce Suspicious or Malicious Attack Score verdicts, `spam` flags produce Spam verdicts, and `graymail` flags produce Graymail verdicts, with ASA automatically reviewing Suspicious messages.

WAS THIS USEFUL?



NETRESEC

Feedback



CapLoader 2.2 Released (2026-10-19)

BLOG ↗

RANK

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

CapLoader 2.2 expands its alerting engine with abuse-risk detections and clearer alert-based visualization, alongside usability improvements to flow transcripts and BPF filter authoring.

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

01 Abuse-risk alerting and visualization in CapLoader

NEW

33

CapLoader 2.2 adds built-in detections for services and infrastructure that may be abused, expanding CapLoader's alerting engine, and improves alert-based coloring for traffic visualization.

- Names the capability but no detection list or mechanism given [product docs](#)

02 Smarter BPF filter feedback

IMPROVED

28

CapLoader 2.2 adds more helpful BPF filter feedback to guide filter authoring.

- Names BPF filters but no example of feedback shown [product docs](#)

03 Improved flow transcript navigation

IMPROVED

20

CapLoader 2.2 improves flow transcripts to be easier to navigate.

- Brief statement with no detail on what changed [product docs](#)

COVERAGE

NEW detection [Abuse-focused infrastructure alert](#)

NEW integration [External blocklist support](#)

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

CapLoader 2.2 Released (2026-10-19)

BLOG ↗

LATEST

STABLE

CapLoader 2.2 adds an AlertSeverity color scheme, arrow-key transcript navigation with pre-fetching, Enter-to-open-transcript from any tab, and BPF filter color/tooltip error feedback.

- › Adds the **AlertSeverity** row color scheme, coloring each flow, service, and host by the highest-severity alert associated with it.
- › Enables keyboard navigation between flow transcripts using the arrow keys, with pre-fetching of the next Flow Index for faster loading.
- › Pressing Enter on any tab now opens the transcript for the selected flow(s), not only from the Flows tab.
- › Adds color and tooltip feedback in BPF filter boxes to warn about possible syntax mistakes.

WAS THIS USEFUL?



CLOUD & CONTAINER SECURITY

Upwind

Feedback



Runtime First, Agentless Everywhere (2026-10-19)

BLOG ↗

RANK

Upwind is a runtime-powered CNAPP that protects cloud workloads by combining posture, detection and response, and identity security with real-time runtime context.

Upwind shipped a Kubernetes Agentless Scanner with a three-layer security model for clusters that can't run a node sensor, added runtime detection for AI-agent activity, and launched a new Databricks Security integration covering Unity Catalog inventory, posture and runtime-ranked risk.

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

01 Databricks Security integration for Lakehouse and AI assets

NEW

90

Adds Databricks Security via **Components > Integrations**, covering asset inventory for workspaces, metastores, catalogs, schemas, grants, storage credentials, external locations, Delta Sharing shares and recipients, model serving endpoints, vector search indexes, and registered models. It connects Unity Catalog grants to underlying cloud resources — mapping storage credentials to IAM roles, clusters to compute resources, network configuration to VPCs and security groups, and encryption configuration to KMS keys — in a single traversable graph, surfaces Databricks-specific findings (overly broad data grants, non-expiring access tokens, publicly exposed workspaces and clusters, universally readable secrets, long-lived Delta Sharing credentials) in the same severity queue, framework reports, and workflows as existing cloud findings, and ranks risk by runtime exposure so live model serving endpoints taking traffic are prioritized over registered-but-idle ones.

Connect a Databricks account to Upwind so that Unity Catalog grants, storage credentials, and AI assets appear in your unified cloud security inventory.

📌 In the Upwind console, go to Components > Integrations, add your Databricks account, grant Upwind read access, then choose the workspaces and metastores in scope to begin inventory collection.

– Exact nav path plus full asset and finding list with setup steps

Upwind Launches Databricks Security for...

02 Kubernetes Agentless Scanner for inventory and vulnerabilities

NEW

70

Adds a Kubernetes Agentless Scanner that connects to a cluster without a DaemonSet or node install, builds a full inventory of Deployments, DaemonSets, StatefulSets and services, extracts SBOMs from deployed images, and surfaces vulnerabilities on the workload running each image. Agentless image scans also detect exposed secrets and malicious files in image contents, surfaced in Threats alongside other detections, and the resulting inventory appears under dedicated resource kinds in Inventory, covering both public and private clusters from a single view.



– Explains mechanism and scope but gives no setup command

Runtime First, Agentless Everywhere (20...

03 **Three-layer Kubernetes security architecture** NEW 60

Introduces a three-layer security model — Upwind Sensor (node-level runtime), Kubernetes Agentless Scanner (inside the cluster), and Cloud Scanners (outside the cluster) — all feeding the same inventory, security graph, and Containers & Kubernetes Focus Mode.

– Names the three layers and shared surfaces, no config steps

Runtime First, Agentless Everywhere (20...

04 **Cloud Scanners discover external Kubernetes footprint** NEW 60

Cloud Scanners now surface the Kubernetes cluster's external footprint — managed cluster resources, node VMs and their OSes, load balancers, IAM roles, and network paths — enabling discovery of unknown clusters and converting them into agentless scanner candidates.

– Describes scope and outcome but not how to trigger it

Runtime First, Agentless Everywhere (20...

Runtime detection identifies AI-agent activity — unexpected shells, previously unseen processes, cloud metadata service calls, and uncharted service-to-service traffic — that manifests only at runtime and is invisible to manifest or image scans.

– Lists detection signals but no UI or config entry point

Runtime First, Agentless Everywhere (20...

WAS THIS USEFUL?



◆ Container & Image Security

Anchore Grype

Feedback

i

v0.121.0

NOTES ↗

CODE ↗

RANK

A vulnerability scanner for container images and filesystems

Grype v0.121.0 improves end-of-life detection by falling back to an alternate OS release date field.

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

01

OS SUPPORT_END fallback for EOL warnings

IMPROVED

55

When generating end-of-life warnings, Grype now falls back to the **SUPPORT_END** date from OS release information, so distros that publish a support-end date but no explicit EOL date are correctly flagged.

– Names the exact field used but no usage example or command

v0.121.0

WAS THIS USEFUL?



Chainguard

[Feedback](#)[Week of 2026-10-19](#)[NOTES ↗](#)[RANK](#)

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

Chainguard shipped per-tag image customization via Custom Assembly Overlays, new Beta products covering .NET library management, AI-driven vulnerability remediation, hardened CI and agent isolation, 45 new container images, and announced a mandatory Sigstore v3 signing migration with a December 2026 cutover.

←→ **WHAT SHIPPED · 7 FEATURES** most completely described first [? what's the number?](#)

01 Custom Assembly Overlays for per-tag image customization NEW

90

Custom Assembly Overlays (Early Access) let you apply packages, environment variables, annotations, certificates, and user accounts to specific image tags, variants, or entire organizations via `chainctl` or Terraform, without affecting unrelated versions. The new `image_overlay` Terraform resource manages overlays and bindings, and version templates such as `py{{major}}.{{minor}}-cryptography` let each tag resolve package names to its own version automatically.

– Names exact Terraform resource and template syntax, lacks CLI example

Week of 2026-10-19

02 Sigstore v3-only signing after December 2026

BREAKING

85

After December 15, 2026, Chainguard will sign artifacts with Sigstore v3 format only; Cosign v2 or v3.1.0 and earlier cannot verify artifacts published after that date, so upgrade to Cosign v3.1.1 or later before the deadline. Custom verification tooling built against the Cosign v2 libraries will break because the Cosign library API changed in v3, and registries that do not support the OCI 1.1 Referrers API may fail to copy Sigstore v3 signatures when mirroring Chainguard artifacts after the cutover.

– Exact versions, dates, and API requirement give clear migration path

Week of 2026-10-19

03 45 new container images added NEW

70

Adds 45 new container images this week, including `crossplane-function-kcl`, `crossplane-gcp-bigquery`, `crossplane-gcp-cloudfunctions`, `crossplane-gcp-cloudrun`, `crossplane-gcp-cloudscheduler`, `crossplane-gcp-datastream`, `crossplane-gcp-secretmanager`, `crossplane-azuread`, `crossplane-provider-nop`, multiple `openstack-*` variants, `logstash-fips`, `mountpoint-s3-csi-driver-fips`, `opentelemetry-dotnet-instrumentation`,

`commercial-nginx-ingress-plus-fips`, `openscap-fips`, `nemo-gym`, `nemo-skills`, and others — many with FIPS variants.

– Names many concrete images that can be pulled directly Week of 2026-10-19

04 **Chainguard Libraries for .NET (Beta)** NEW 65

Chainguard Libraries for .NET (Beta) replaces public NuGet feeds with a secure Chainguard endpoint for C#, F#, and Visual Basic dependencies, applying automated malware and greyware scanning plus customizable organization policies.

– Describes mechanism but no endpoint or setup command given Week of 2026-10-19

THINNER COVERAGE BELOW

05 **Chainguard Checks: hardened CI runners** NEW 55

Chainguard Checks (Beta) replaces standard CI runners with hardened microVMs that secure credentials and enforce egress policies by default, usable before code submission so pull requests pass on the first try.

– Describes behavior and workflow but no config or command named Week of 2026-10-19

06 **Chainguard Workspaces: isolated microVMs for AI agents** NEW 55

Chainguard Workspaces (Beta) gives each AI coding agent an isolated, hardened microVM with no access to developer credentials, preventing prompt injection and credential exposure from agents running on developer machines.

– Clear mechanism and benefit, but no setup or API surface named Week of 2026-10-19

07 **Chainguard Harness: AI-driven vulnerability discovery and fixes** NEW 45

Chainguard Harness (Beta) is an AI-powered engine that finds exploitable runtime vulnerabilities — including logic flaws, injection, and resource exhaustion — and returns engineer-reviewed fixes, without requiring customer access to frontier models.

– Explains purpose but no interface, command, or scope detail Week of 2026-10-19

↳ **ALSO FROM THESE RELEASES**

.
. .
.

↳ **BREAKING ON UPGRADE**

! After December 15, 2026, Chainguard will sign artifacts with Sigstore v3 format only; Cosign v2 or v3.1.0 and earlier cannot verify artifacts published after that date — upgrade to Cosign v3.1.1 or later before that deadline.

- ! Custom verification tooling built against the Cosign v2 libraries will break because the Cosign library API changed in v3 — review and update against the v3 libraries.
- ! Registries that do not support the OCI 1.1 Referrers API may fail to copy Sigstore v3 signatures when mirroring Chainguard artifacts after December 15, 2026 — confirm registry support before the cutover.

WAS THIS USEFUL?



DETECTION & RESPONSE

Google SecOps

[Feedback](#)[changelog · Oct 20, 2026](#)[CHANGELOG ↗](#)**RANK**

Google secops is a comprehensive documentation, guides, and resources for {% dynamic print site_values.cloud_name %} products and services.

Google SecOps updated parsers for four log source types and migrated Applied Threat Intelligence curated detection rules to the new **GTI_IOC** feed, a breaking change for custom rules still referencing legacy feeds.

↳ **WHAT SHIPPED · 2 FEATURES** most completely described first [? what's the number?](#)

01 **ATI curated rules migrated to GTI_IOC feed**

BREAKING

65

Applied Threat Intelligence (ATI) curated detection rules now match against the **GTI_IOC** feed, replacing the legacy feeds. Custom YARA-L rules that still reference the legacy ATI feeds must be manually updated to use **GTI_IOC**, or they will no longer match on the curated indicator data.

:

– Names feed, migration requirement, and consequence of inaction

[changelog-20261020-c58095ea](#)

THINNER COVERAGE BELOW

02 **Parser updates for four log types**

IMPROVED

40

Updates parsers for **FORTINET_SANDBOX**, **PURE_STORAGE**, **RSA_NETWITNESS**, and **UPTYCS_EDR** log types.

– Names log types but no detail on what changed in parsing

[changelog-20261020-c58095ea](#)

↳ **BREAKING ON UPGRADE**

! Custom YARA-L rules that still reference the legacy ATI feeds must be updated to use the **GTI_IOC** feed, or they will no longer match on the curated indicator data.

WAS THIS USEFUL?



Tines

Feedback



changelog · Oct 20, 2026

CHANGELOG

RANK

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

Tines' 43.7.0 release adds a faster Records API v2, MCP-based CRUD and search for Records, Claude Sonnet 5.5 as a model option, and SSO group app restrictions, alongside several Cases UI improvements and advance notice of PostgreSQL 14 and Docker Hub deprecations.

WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 Docker Hub image deprecation for self-hosted

BREAKING

75

New Tines images will only be released to `oci.tines.com` beginning January 15, 2027, with no new images pushed to Docker Hub after that date; access to Tines images on Docker Hub ends entirely on March 1, 2027, so Docker deployments must migrate to the OCI registry before then.

– Exact registry and both deadlines given, no migration command

changelog-20261020-f89eac3e

02 PostgreSQL 14 end-of-life migration

BREAKING

70

PostgreSQL 14 reaches end of life on November 12, 2026, so self-hosted deployments must migrate to PostgreSQL 18 before then.

– Clear deadline and target version but no migration steps

changelog-20261020-f89eac3e

THINNER COVERAGE BELOW

03 Records API v2

NEW

55

Adds `Records API v2`, which requires fewer calls per read and performs faster searches than the previous version.

– Names the API version but no endpoints or methods

changelog-20261020-f89eac3e

04 MCP support for Records

NEW

55

Adds Tines MCP support for Records, enabling creating, updating, and deleting Records as well as discovering, reading, and searching Records.

– Lists supported operations but no commands shown


changelog-20261020-f89eac3e



05 Cases UI improvements


IMPROVED

55

Several Cases improvements: sets background colors for file and group blocks, downloads all files in a case block at once, shows when a case block was last updated, and adds Tines 3B workflows as case actions.







– Enumerates all four constituent changes with no mechanism depth

[changelog-20261020-f89eac3e](#)

06 Claude Sonnet 5.5 model support NEW 35

Adds Claude Sonnet 5.5 as an available model in Tines.

– Bare mention, no usage detail [changelog-20261020-f89eac3e](#)

07 SSO group app restrictions NEW 35

Restricts apps to specific SSO groups.

– States capability but no config mechanism named [changelog-20261020-f89eac3e](#)

08 Self-hosted database overview page NEW 35

Adds a self-hosted database overview page.

– Names a UI page but no detail on contents [changelog-20261020-f89eac3e](#)

↳ **BREAKING ON UPGRADE**

- ! PostgreSQL 14 reaches end of life on November 12, 2026 — self-hosted deployments must migrate to PostgreSQL 18.
- ! New Tines images will only be released to oci.tines.com beginning January 15, 2027 — no new images will be pushed to Docker Hub after that date.
- ! Access to Tines images on Docker Hub ends March 1, 2027 — Docker deployments must migrate to the OCI registry before this date.

WAS THIS USEFUL?



Watcher

[Feedback](#)

v3.7.0

[NOTES](#)[RANK](#)

Watcher - Open Source AI-powered Cyber Threat Intelligence & Hunting Platform. Developed with Django & React JS.

Watcher v3.7.0 introduces bulk create/update/delete REST endpoints across every module alongside matching bulk-focused UI, plus dedicated endpoints to export suspicious domains to Legitimate Domains or MISP.

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

01 Bulk create/update/delete endpoints across all modules NEW

88

Adds `POST <endpoint>/bulk_create/`, `bulk_update/`, and `bulk_delete/` endpoints to every module's API, each returning a per-item report of `created` / `skipped` / `errors`. Request size is capped by new `BULK_CREATE_MAX_ITEMS` and `BULK_ACTION_MAX_ITEMS` env vars in `.env`, both defaulting to `200`, with requests above the limit rejected wholesale.

Raise the bulk-create ceiling to 500 lines per request when ingesting large domain threat feeds into Watcher.

```
ini

BULK_CREATE_MAX_ITEMS=500
BULK_ACTION_MAX_ITEMS=500
```

— Names exact endpoints, env vars, defaults and a config example. v3.7.0

02 Bulk export of domains to Legitimate Domains or MISP NEW

81

Adds `bulk_export_legitimate/` and `bulk_export_misp/` endpoints on `/api/site_monitoring/site/` to export selected suspicious domains to Legitimate Domains or MISP, generating one MISP event per site.

— Exact endpoint path and behaviour given, no usage example. v3.7.0

03 Bulk-add and bulk-action UI across modules NEW

61

Adds a `One item per line` bulk-add toggle to add forms for Legitimate Domains, Suspicious Domains, DNS Threats Monitored (domains and keywords), Data Leak search patterns, Banned Words, and Monitored Keywords, plus a Preview button on bulk-add forms to show what would be created without saving. Tables across the app gain bulk row actions so ticking multiple rows can change legitimacy, takedown/legal-

team/blocking requests, alert level, or status in one step, or delete them together; lists and dashboard statistics now refresh without a full page reload after these bulk actions.

– Describes UI surfaces and modules but gives no exact paths. v3.7.0

WAS THIS USEFUL?  

OpenCTI

Feedback 

7.261019.0

NOTES 

RANK

Open Cyber Threat Intelligence Platform

OpenCTI 7.261019.0 adds dashboard variables for dynamic widget filtering and proxies XTM One chatbot content.

 **WHAT SHIPPED** • **2 FEATURES** most completely described first [? what's the number?](#)

01 **Dashboard variables management UI** NEW 40

Adds a dashboard variables management UI and support for using variables in widget filters, enabling dynamic dashboard customization.

– Names UI and widget filter usage but no mechanism details. 7.261019.0

02 **XTM One chatbot content proxying** NEW 40

Proxies decks, reports, and videos from XTM One into the embedded chatbot interface.

– Names content types and source but not how to configure it. 7.261019.0

WAS THIS USEFUL?  

Cloudmersive is a developer API platform offering cloud-based services for file scanning and malware detection, document conversion, image processing, and data validation.

Cloudmersive's Fraud Detection AI API gains two new input parameters for finer-grained fraud analysis.

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

01 Country code and document class parameters in Fraud Detection AI

NEW 45

The Fraud Detection AI API now accepts `country code` and `document class` input parameters, released in Cloud Fraud AI API v2-1-0.

– Names exact parameters but no usage mechanism or example given

Cloudmersive Cloud Fraud AI API v2-1-0 ...

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

rsigma

[Feedback](#)[v0.25.0](#)[NOTES ↗](#)[RANK](#)

A complete Sigma detection engineering toolkit: parser, linter, evaluator, correlation engine, conversion framework, streaming daemon, MCP and LSP servers

rsigma v0.25.0 ships a wave of breaking changes tightening daemon/MCP security, HIR caching, and core runtime types, alongside new inhibit-source capping, `engine eval -o` file output, custom logsource extractors for embedders, and an OTLP decode fuzz harness.

↳ WHAT SHIPPED • 16 FEATURES most completely described first ? what's the number?

01 HIR retention control and cache versioning

BREAKING 90

Adds an opt-in `Engine::set_retain_hir(true)` method (must be called before rules are loaded) to control whether loaded rule HIR is retained; `save_hir` now returns `EvalError::HirCache` unless retention was enabled. `HIR_SCHEMA_VERSION` is bumped to 4, making HIR cache files from prior versions incompatible.

– Names method, error variant, and schema version with clear before/after [v0.25.0](#)

02 Inhibit source capping in alert pipeline

NEW 85

Adds `max_inhibit_sources` top-level key (default 100000) to alert pipeline config to cap tracked inhibit sources, exposing a new `rsigma_inhibit_sources_refused_total{rule}` counter plus `AlertPipelineFile.max_inhibit_sources` and `MetricsHook.on_inhibit_source_refused` hooks for library users.

– Names config key, counter, and library hooks but no worked example [v0.25.0](#)

03 Secure-by-default daemon and MCP HTTP server

BREAKING 85

Non-loopback daemon API access now requires authentication or the explicit `--allow-unauthenticated` flag, and `mcp serve --http` now requires either `--rules-dir` or `--allow-unconfined-paths`; existing unauthenticated or unconfigured deployments will be refused until updated.

– Two exact flags named, runnable migration path [v0.25.0](#)

04 `EventFilter` now takes events by value

BREAKING 85

`rsigma_runtime::EventFilter` is now `dyn Fn(serde_json::Value) -> Vec<serde_json::Value>`, taking the parsed line by value; custom filters returning `vec![v.clone()]` must be updated to return `vec![v]`.

– Exact type signature and migration fix spelled out v0.25.0

05 Shared ``${source.*}`` reference grammar IMPROVED 70

Adds `rsigma_eval::pipeline::sources::scan_source_templates` as the single shared grammar for ``${source.*}`` references, now accepting ids containing digits, underscores, and hyphens; a reference id containing other characters (e.g. a dot or space) now fails to load with an error naming the id.

– Names function and exact validation rules with before/after v0.25.0

06 `'SigmaCollection::diagnostics'` replaces `'errors'`

BREAKING

65

`SigmaCollection::errors` is replaced by `SigmaCollection::diagnostics`, separating parse warnings from parse errors; code calling `.errors()` will not compile.

– Exact API rename with compile-time consequence v0.25.0

07 File output for `'engine eval'` NEW

60

Adds `engine eval -o` to write evaluation results to a file instead of stdout.

– Exact runnable flag, but no further detail on format v0.25.0

08 OTLP decode fuzz harness NEW

60

Adds `fuzz_otlp_logs_decode` cargo-fuzz target decoding arbitrary bytes as `ExportLogsServiceRequest` and asserting valid JSON output, run in a weekly fuzz workflow from a committed seed corpus.

– Internal testing infrastructure, not user-facing action v0.25.0

09 `'EvalError::Pipeline'` variant added

BREAKING

60

Adds an `EvalError::Pipeline` variant and marks `EvalError` as `# [non_exhaustive]`; exhaustive match arms on `EvalError` will not compile.

– Named variant and compile-time impact but thin rationale v0.25.0

THINNER COVERAGE BELOW

10 `'incompatible_modifiers'` promoted to lint error

BREAKING

55

`incompatible_modifiers` is now a lint error rather than a warning; pipelines with incompatible modifiers will fail to load.

– Named lint with clear before/after severity v0.25.0

11 Zero-copy JSON event value access NEW 50

Adds `JsonEvent::into_value` method to unwrap a JSON event value without cloning when the event owns it.

– Named API with clear purpose but no example v0.25.0

12 Per-component URL encoding in enricher/webhook templates 50

BREAKING

Event values in enricher and webhook URLs are now encoded as one component each; templates relying on unencoded values will produce different output.

– Behavior change described but no migration example v0.25.0

13 Dead-letter queue write failure counter NEW 45

Adds `rsigma_dlq_write_failures_total` counter to track failed writes to a configured DLQ sink.

– Named metric but no mechanism or usage detail v0.25.0

14 Unified `NumericLiteral` numeric type 45

BREAKING

Numeric matchers and backend methods now use `NumericLiteral`; code using the prior numeric type will not compile.

– Named type but no detail on prior type or migration v0.25.0

15 `'cel'` 0.15 upgrade changes expression evaluation 40

BREAKING

Upgrading to `cel` 0.15 changes how some extract expressions evaluate; existing `CEL` expressions should be reviewed against the migration notes.

– Dependency bump noted but no specifics of what changed v0.25.0

16 Pluggable custom logsource extractors NEW 35

Supports pluggable custom logsource extractors for embedders via a new embedder API.

– Named capability but no API surface detail v0.25.0

↳ **BREAKING ON UPGRADE**

! A non-loopback daemon API now requires authentication or the explicit `--allow-unauthenticated` flag; existing unauthenticated non-loopback setups will be refused.

- ! `mcp serve --http` now requires `--rules-dir` or `--allow-unconfined-paths` ; existing invocations without these flags will be refused.
- ! Event values in enricher and webhook URLs are now encoded as one component each; templates relying on unencoded values will produce different output.
- ! `HIR_SCHEMA_VERSION` is now 4; existing HIR cache files from prior versions are incompatible.
- ! `save_hir` now returns `EvalError::HirCache` unless `Engine::set_retain_hir(true)` was called before loading rules; callers that relied on implicit HIR retention will now get an error.
- ! `SigmaCollection::diagnostics` replaces `SigmaCollection::errors` ; code calling `.errors()` will not compile.
- ! `EvalError` gains the `Pipeline` variant and is now `#[non_exhaustive]` ; exhaustive match arms on `EvalError` will not compile.
- ! `rsigma_runtime::EventFilter` is now `dyn Fn(serde_json::Value) -> Vec<serde_json::Value>` , taking the parsed line by value; custom filters returning `vec![v.clone()]` must be updated to return `vec![v]` .
- ! Numeric matchers and backend methods now use `NumericLiteral` ; code using the prior numeric type will not compile.
- ! `incompatible_modifiers` is now a lint error rather than a warning; pipelines with incompatible modifiers will fail to load.
- ! `cel` 0.15 changes how some extract expressions evaluate; existing CEL expressions should be reviewed against the migration notes.
- ! A `${source.*}` reference id containing characters other than ASCII letters, digits, `_` , or `-` (e.g. a dot or a space) now fails to load with an error naming the id.

WAS THIS USEFUL?



EXPOSURE & VULNERABILITY MANAGEMENT

Rapid7 InsightVM

Feedback



changelog · Oct 20, 2026

DOCS ↗

BLOG ↗

Rapid7 InsightVM is a vulnerability management platform that scans networks and assets to identify, prioritize, and help remediate security weaknesses.

Bulk Export API now honors any API key's access level; site alerts restricted to users with site access; FIPS-mode scan engines reconnect automatically after console restart.

- › Bulk Export API now accepts any API key and honors that key's access level, removing the previous requirement for an organization-level or platform administrator key.
- › Stronger access controls for site alerts make alert settings visible only to users with access to the site the alert belongs to.
- › Reverse-paired scan engines (engine-to-console pairing) now reconnect automatically after a Security Console restart in FIPS mode, with no new shared secret and no engine update required.

COVERAGE

NEW

compliance framework [CIS Ubuntu 26.04 Benchmark v1.0.0 policy content](#)

[All of Rapid7 InsightVM's coverage changes →](#)

WAS THIS USEFUL?



DefectDojo Pro 3.4.200 ships a major Findings grid and filter overhaul, dedupe-match automation, new Asset relationship types with exposure propagation, in-app walkthroughs, and new StackHawk and Qualys TotalAppSec connectors, alongside breaking changes to license enforcement and the sidebar.

←→ WHAT SHIPPED · 13 FEATURES most completely described first ? what's the number?

01 Findings list grid view and filter bar NEW 85

The Findings list gains a grid view and filter bar — also applied to the Assets, Organizations, Engagements, and Tests lists — with a Match Any mode, an Owner (Team) column and filter, API support for filtering Findings by owner and assignee, In List and Not in List modes for text column filters, an Include / Exclude switch on every multi-select column filter, and paste-a-list support for Asset and Organization pick-list filters.

— Names many concrete UI/API mechanics but no exact navigation path given product docs

02 Faster fresh installs via schema pre-loading IMPROVED 80

Fresh installs now load a per-release schema snapshot into empty databases on first start, speeding up deployment. The `DD_SCHEMA_DUMP_PATH` config key lets operators override the schema file path, and setting it to an empty value disables schema pre-load entirely.

Prevent a fresh-install from running every migration by pointing to a pre-built schema dump, or disable the feature entirely.

```
env  
  
DD_SCHEMA_DUMP_PATH=/opt/defectdojo/schema.sql
```

— Named config key with runnable example and clear behaviour 3.4.200

03 Zero-downtime mcp-server rollout wait config NEW 75

The `lifecycle.preStop.waitSeconds.mcpServer` config key sets the preStop wait during zero-downtime `mcp-server` rollouts, so in-flight MCP requests can complete before the pod terminates.

Control zero-downtime mcp-server rollout drain time so in-flight MCP requests complete before the pod terminates.

```
yaml

lifecycle:
  preStop:
    waitSeconds:
      mcpServer: 30
```

– Exact config key with runnable YAML example [product docs](#)

04 **Reachability probe settings for superusers** NEW 67

Superusers can configure a reachability probe under **Settings > Reachability Probe**, which checks addresses held in Locations.

Configure the reachability probe as a superuser to check network addresses recorded in Locations.

-  1. Log in as a superuser. 2. Navigate to Settings > Reachability Probe. 3. Enter the addresses to probe and save.

– Exact UI path plus step-by-step usage example [product docs](#)

05 **Dedupe match triggers and batch actions** NEW 65

An On Dedupe Match trigger lets rules mark a matched Finding a duplicate or keep both active with independent SLAs and a link; Copy Status from Matched Finding and Apply Status to Linked Findings actions propagate false positive, out of scope, and risk accepted decisions across linked Findings. A batch dedupe match action is now caller-configurable via the API, giving API consumers control over deduplication behaviour.

– Mechanism described well but no endpoint or exact API call named [3.4.200](#)

06 **License required for database migration** BREAKING 65

The initializer now refuses to migrate without a valid license, requiring one via **DD_LICENSE**, the mounted license file, or the database; operators should consult the upgrade guide before upgrading.

– Names the env var but no concrete migration command shown [product docs](#)

07 **Asset relationship types and deploys_to exposure propagation** NEW 64

DefectDojo Pro ships four built-in Asset relationship types by default, including **deploys_to**, which propagates exposure from host clusters to deployed services — a service deployed onto an internet-facing cluster inherits the cluster's effective exposure. **deploys_to** is intentionally excluded from Finding roll-up so a cluster's Findings are not attributed to every service it hosts.

⋮

– Explains mechanism and scope but no setup steps given [product docs](#)

08 In-app walkthroughs in the Pro UI NEW 60

The Pro UI adds in-app walkthroughs — an App Tour plus task walkthroughs for importing, triage, access, reports, admin setup, the Jira connector, and the Triage Engine — controlled by the **In-App Walkthroughs** feature flag, which is on by default.

– Names the feature flag and walkthrough scopes but no navigation steps [product docs](#)

THINNER COVERAGE BELOW

09 Asset type automation via Triage Engine and connectors IMPROVED 55

Triage Engine rules can now set an Asset's type and use Asset metadata, custom fields, and contacts in notifications, with select fields supporting partial matching; connectors also set Asset type from vendor data on every sync, without overwriting a type set by a person or rule.

– Behaviour explained but no config key or UI path named [product docs](#)

10 Connector enhancements for Wiz, FOSSA, and Qualys IMPROVED 55

Wiz can report network exposure as asset exposure evidence; FOSSA syncs can be limited by issue category and by project label or title; Qualys can import host tags and group Assets by tag.

– Names three connectors and their new sync capabilities, no setup steps [product docs](#)

11 Classic sidebar removed BREAKING 33

The classic sidebar and the Menu 2.0 toggle are removed; all users are migrated to the current sidebar.

– Bare statement of removal with no further detail [product docs](#)

12 Risk meter on Findings NEW 33

Findings now show a risk meter that marks the Finding's band visually instead of displaying a written level.

– Brief description with no further mechanism or config [product docs](#)

13 New StackHawk and Qualys TotalAppSec connectors NEW 30

DefectDojo Pro adds StackHawk and Qualys TotalAppSec connectors.

– Only named, no configuration or capability detail given [product docs](#)

L--> ALSO FROM THESE RELEASES

.
. .
.

L--> BREAKING ON UPGRADE

- ! The initializer refuses to migrate without a valid license in `DD_LICENSE`, the mounted license file, or the database — see the upgrade guide before upgrading.
- ! The classic sidebar and Menu 2.0 toggle are removed; all users are migrated to the current sidebar.

WAS THIS USEFUL?



faraday-cli

[Feedback](#)[2.3.0](#)[NOTES](#)[RANK](#)

Faraday's Command Line Interface

faraday-cli 2.3.0 adds a full FaradAI cloud-agent surface — scanning, triage, and SaaS onboarding — alongside new multi-session authentication management, letting users scan without a self-hosted Faraday server.

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

01 Multi-session auth management

NEW

95

Adds `auth list` to display all saved connection sessions with columns ID, ACTIVE, NAME, KIND, SERVER, USER, WORKSPACE and mark the active one; `auth use <id|name|user>` to switch the active connection by ID, name, or user/email; `auth rm <id|name|user>` to remove a saved session (clearing the active connection if it was active); and a `-n / --name` flag on login to set a custom session name instead of defaulting to the server host.

Switch between two analyst accounts on the same Faraday server without logging out — list sessions then activate the one you need.

```
$ faraday-cli auth list
faraday-cli auth use bob@corp.faradaysec.com
```

Switch between named Faraday connection sessions when working across multiple environments.

```
$ faraday-cli auth list
faraday-cli auth use prod-server
```

— Names every subcommand, flag and column verbatim with runnable examples [2.3.0](#)

02 Cloud agent scanner management

NEW

95

Adds `faraday-cli scanners` command group: `scanners list` enumerates available cloud agents with ID, NAME, SLUG, CATEGORY, `LAST_RUN` columns plus `-j / --json-output / --pretty` flags; `scanners get <scanner_id>` shows a single agent and the parameters it accepts; `scanners run` launches an agent against one or more workspaces using `--agent / -a`, `--workspace`, `--param KEY=VALUE`, `--params-json` (to pass all parameters as one JSON object), and `--watch` (polls every ~15s, printing each finished command).

Run a specific named cloud agent with raw parameters when you need to target a non-default FaradAI scanner.

```
$ faraday-cli scanners run --agent <id|name>
```

Run a specific cloud agent by name or ID with custom raw parameters.

```
$ scanners run --agent <id|name>
```

– All subcommands, flags and columns named verbatim 2.3.0

03 FaradAI pentest and triage commands NEW

95

Adds `faraday-cli pentest` (shortcut for `faradai scan`) to launch an AI-driven pentest against a target via the FaradAI cloud agent, using `--target` for the URL, `--watch` to tail the run to completion, and `--advanced` for a deeper 80-round audit; and `faraday-cli faradai triage` to triage findings in a workspace, supporting flags such as `--severities`, `--confirmed-only`, `--auto-close-fp`, `--reproduce`, and `--dry-run` to preview changes before committing.

Run an advanced FaradAI pentest (80 rounds, full-audit instruction) against a target and stream results to the terminal until it completes.

```
$ faraday-cli pentest --target https://example.com --workspace my_workspace --advanced --watch
```

Triage only critical and high confirmed findings in a workspace, auto-closing false positives, without committing any changes — then re-run without `--dry-run` when satisfied.

```
$ faraday-cli triage --workspace my_workspace --severities critical,high --confirmed-only --auto-close-fp --reproduce --dry-run
```

Run an AI-driven pentest against a web target and block until it finishes — useful in a CI gate where you need results before proceeding.

```
$ faraday-cli pentest --target https://example.com --watch
```

Trigger the deeper 80-round audit mode for a more thorough AI pentest engagement.

```
$ faraday-cli pentest --target https://example.com --advanced --watch
```

Launch an AI pentest against a web target and stream results to your terminal until the scan completes.

```
$ pentest --target https://example.com --watch
```

Run a thorough 80-round AI audit against a target when you need deeper coverage than the default scan.

```
$ pentest --target https://example.com --advanced --watch
```

– Named flags and runnable commands for both pentest and triage 2.3.0

04 FaradAI SaaS onboarding commands NEW

70

Adds `faraday-cli faradai register`, `faradai verify`, and `faradai login` subcommands to sign up for and authenticate against the FaradAI SaaS service, enabling scanning without a self-hosted Faraday server.

Register and verify a new FaradAI SaaS account entirely from the CLI to start scanning immediately.

```
$ faraday-cli faradai register
faraday-cli faradai verify
faraday-cli faradai login
```

– Names the three subcommands but no further mechanism detail 2.3.0

WAS THIS USEFUL?



Censys

Feedback



October 19, 2026 (2026-10-19)

BLOG

RANK

Censys is a internet-wide scanning and search platform that discovers and monitors hosts, services, and certificates to help security teams assess their exposure.

Censys shipped a new Threat Hunting MCP Server for AI agent access to its data, alongside a reorganized Platform UI, an in-app feedback portal, 10 new software fingerprints, hundreds of new ASM CVE risks, and five new ARC Rapid Response advisories.

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

01 Platform UI navigation reorganization IMPROVED 62

Reorganizes the Censys Platform UI and search actions into three sections: Investigate, Defend, and Intelligence. Moves search history to its own tab on the upper right corner of the search results page, renames the Reports tab to Aggregate , and relocates documentation, Academy, and Data Definitions resources to the Help and Feedback menu in the top right.

.
. .
. .

- Names exact UI locations and renamed elements October 19, 2026 (2026-10-19)

THINNER COVERAGE BELOW

02 10 new software fingerprints NEW 55

Adds 10 new software fingerprints to the Platform: Ahsay Cloud Backup Suite (AhsayCBS), IBM MCP Context Forge , Nextcloud Collabora Online Built-in CODE Server , OliveTin , Perforce P4 Search (formerly HelixSearch), phpBB , pyLoad , Softaculous Virtualizor , ThinkPHP , and ZITADEL .

- Fully enumerates named fingerprints but no usage detail October 19, 2026 (2026-10-19)

03 Five new ARC Rapid Response advisories NEW 50

Publishes five new Censys ARC Rapid Response advisories covering AnyDesk Linux , GitLab CI/CD , Ivanti Connect Secure , Collabora Online Built-in CODE Server , and BeyondTrust Remote Support/Privileged Remote Access .

- Names all five advisory topics but not content October 19, 2026 (2026-10-19)

04 Threat Hunting MCP Server NEW 40

Introduces the Censys Threat Hunting MCP Server, allowing AI assistants and agents to interact with Censys threat intelligence data through the Model Context Protocol (MCP).

– Names the protocol but no setup or connection details [product docs](#)

05 In-app Platform Feedback Portal NEW 33

Adds an in-app Censys Platform Feedback Portal for submitting feedback and tracking feature requests.

– Describes purpose but not location or workflow [October 19, 2026 \(2026-10-19\)](#)

06 New ASM CVE risks for JS libraries NEW 25

Adds several hundred new CVE risks for JavaScript libraries and web component fingerprints to ASM.

– No specific CVEs or scope named [October 19, 2026 \(2026-10-19\)](#)

WAS THIS USEFUL?  

syzkaller

[Feedback](#)[changes since 2026-10-16](#)[CODE ↗](#)

syzkaller is an unsupervised coverage-guided kernel fuzzer

Records VM environment on dashboard Build entities; syz-agent matches targets by VMType; LTS RC fuzzing sessions are now prioritized in syz-cluster.

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-16
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-16
```

- › Records VM environment (`BuildEnv` , including `VMType` and `MachineType`) on dashboard Build entities, enabling `syz-agent` to match target configurations by `VMType` .
- › Adds `linux/amd64` GCE target to `syz-agent` prod config.
- › Prioritizes LTS RC fuzzing sessions in `syz-cluster` .

COVERAGE

NEW

supported platforms

[Nested KVM state pages mutation for x86, KVM nested state modeled as a resource in sys/linux](#)

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

WAS THIS USEFUL?



◆ Endpoint Management (UEM/RMM)

Iru

[Feedback](#)[What's New in Iru | October 16th](#)[NOTES ↗](#)[LATEST](#)[STABLE](#)

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

Iru adds easier corporate policy access for end users and supports background downloads of custom app installers up to 30 GB.

- › Enables easier access to corporate policies for end users.

COVERAGE

[NEW](#)

supported platform [Windows 11 26H2 managed upgrade support](#)

[NEW](#)

integration [Auto Apps milestone — 500 Auto Apps](#)

[All of Iru's coverage changes →](#)

▼ 1 more release in this issue · 2026-10-09 → 2026-10-16

[What's New in Iru | October 9th](#)[NOTES ↗](#)[STABLE](#)

Iru adds the Auditor Engagement Hub for auditor collaboration in Compliance Automation, new Mac restrictions, and a centralized Identity activity view.

- › Adds the Auditor Engagement Hub inside Compliance Automation to enable direct collaboration with external auditors.
- › Introduces new Mac device restrictions and increased Mac device visibility for improved endpoint control.
- › Adds a centralized view for all Identity activities.

WAS THIS USEFUL?



IDENTITY, ACCESS & DATA PROTECTION

Aembit

Feedback



changelog · Oct 20, 2026

CHANGELOG

RANK

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's biggest MCP-focused release yet adds session visibility and revocation, schema-validated tool outputs, a change-notification relay, and expanded authentication options for MCP Server Workloads, while bringing the Python and TypeScript Edge SDKs to parity across nine Trust Providers — alongside two breaking changes to MCP Identity Gateway error payloads and token expiry handling.

↳ WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 Trust Provider parity across Edge SDKs NEW 90

Adds AWS Metadata Service, Azure Instance Metadata Service, GCP Identity Token, Kubernetes Service Account, and OIDC ID Token Trust Provider support to the Python Edge SDK (`pip install aembit-edge-sdk` , requires Python 3.10+), and adds GitLab, Kubernetes Service Account, and Terraform Cloud Trust Provider support to the TypeScript Edge SDK, bringing both SDKs to parity across the same nine Trust Providers.

Install the new Aembit Python Edge SDK on a Python 3.10+ workload to unlock all nine Trust Providers.

```
$ pip install aembit-edge-sdk
```

– Runnable install command plus every named provider `changelog-20261020-b9732860`

02 Change-notification relay in MCP Identity Gateway NEW 85

Adds `subscriptions_workload_refresh_interval` timeout setting (default: 30 seconds) to MCP Identity Gateway, controlling how often it checks for assignment and rule changes, and relays `subscriptions/listen` change notifications to MCP clients using the 2026-07-28 MCP specification revision — including notifications triggered by Server Workload assignment or MCP Tool Access Control rule changes, and on reconnection.

– Named config key with default but no setup steps shown `changelog-20261020-b9732860`

03 Breaking changes to MCP Identity Gateway error handling 85

BREAKING

The `status` field in MCP client error `data` (sent when MCP Identity Gateway cannot connect to an MCP server) is replaced by `upstreamHttpStatus` — update any client code that reads `status`. MCP Identity Gateway also now accepts an access token for up to 3 seconds after expiry, down from 60 seconds, so MCP clients must refresh access tokens before they expire.

– Exact field rename and expiry window with migration note `changelog-20261020-b9732860`

04 Sessions tab for MCP identity provider sessions NEW

75

Adds `Sessions` tab to each Identity Provider in the Aembit Tenant UI, listing active MCP user sessions with the ability to filter, revoke individual sessions, or revoke all sessions for a user or MCP client at once.

Revoke all active MCP sessions for a compromised user without deleting the Identity Provider.

 In the Aembit Tenant UI, go to Identity Providers > <your IdP> > Sessions, filter by the user, then select `Revoke all sessions for this user`.

– Exact UI navigation and revoke action given `changelog-20261020-b9732860`

05 Authentication methods for MCP Server Workloads NEW

75

Adds `Header` authentication scheme support to MCP Identity Gateway, sending credentials to MCP servers in a named header so MCP servers that accept an API key via a custom header now receive it correctly. Also extends MCP Server Workloads to support any authentication method available to HTTP Server Workloads, including `API Key`, `HTTP Authentication` with the `Basic` scheme, and `Static HTTP Headers`.

– Names every auth scheme but no config steps shown `changelog-20261020-b9732860`

06 New event logging in MCP Identity Gateway NEW

65

Adds the `access.registration` event type to Access Authorization Events, recording each Dynamic Client Registration (DCR) request an MCP client sends — successful registrations log client ID, client name, redirect URIs, grant types, and response types at severity `Info`, while failed registrations log at severity `Error` with a `reason` and the failing `attribute` (and its supplied value when present); also adds per-MCP-server workload events for `tools/list` and `resources/list` fan-out requests, including events for servers that fail or are skipped because they do not support the method, naming the Server Workload and recording the server's host and HTTP status when known.

– Names exact event fields but no UI path to view them Describes event content but no path to view events

07 Integration guides and Edge SDK logging reference NEW

60

Adds new integration guides for Azure Functions, Azure virtual machines, [GitHub Actions](#), [GitLab CI/CD](#), Kubernetes, and Terraform Cloud, each covering both TypeScript and Python, and adds an Edge SDK logging reference documenting events the SDKs log, how to route them to a logger, and which hosts the SDKs contact.

– Lists guide topics but no direct link or command [changelog-20261020-b9732860](#)

THINNER COVERAGE BELOW

08 Tool-output schema validation in MCP Identity Gateway NEW

50

Adds tool-output validation to MCP Identity Gateway: when a tool publishes an output schema, completed results are checked against it and the MCP client receives an error in place of any result that does not match.

– Explains mechanism but no flag to configure it [changelog-20261020-b9732860](#)

↳ **BREAKING ON UPGRADE**

- ! The `status` field in MCP client error `data` (sent when MCP Identity Gateway cannot connect to an MCP server) is replaced by `upstreamHttpStatus` — update any client code that reads `status`.
- ! MCP Identity Gateway now accepts an access token for up to 3 seconds after expiry, down from 60 seconds — MCP clients must refresh access tokens before they expire.

WAS THIS USEFUL?



Okta for AI Agents

Feedback



changelog · Oct 20, 2026

CHANGELOG

RANK

Okta for AI Agents governs AI agent identities, permissions, and access across autonomous workflows.

Okta's October 2026 release (2026.41.0) extends AI agent governance with imports from Snowflake, Workato, and Cursor, while also shipping NFC badge authentication, PowerShell execution on-premises AD, app-by-app Identity Engine migration, and new fraud and passkey controls across the broader platform.

WHAT SHIPPED · 11 FEATURES most completely described first ? what's the number?

01 Self-service registration fraud policy NEW 75

Adds a self-service registration policy on the [User Profile Policies](#) page for Identity Threat Protection customers, enabling rules to automatically deny sign-ups from disposable email domains or specific network zones.

Block self-service sign-ups from disposable email domains under the ITP self-service registration policy to reduce fake account creation.

In the Admin Console, go to Security > User Profile Policies, open the self-service registration policy, and add a rule targeting disposable/throwaway email domains or specific network zones to set the action to Deny.

– Includes exact admin console path and rule configuration steps

changelog-20261020-2c40ee6d

02 New System Log events for telephony status and certification staging

NEW

70

Adds new System Log events: `system.sms.receive_status` and `system.voice.receive_status` when custom telephony providers report SMS/voice delivery status, and `certification.campaign.stage.start` / `certification.campaign.stage.end` for Access Certification campaign staging, with `outcome.result` reporting `SUCCESS` or `FAILURE`.

– Names all event types and the outcome field verbatim

changelog-20261020-2c40ee6d

THINNER COVERAGE BELOW

03 PowerShell execution in on-premises Active Directory NEW

Enables PowerShell script execution in on-premises Active Directory environments via the Active Directory agent, invocable through Okta Workflows using the Okta public API.

– Names agent, Workflows, and public API as entry points `changelog-20261020-2c40ee6d`

04 App-by-app Identity Engine migration with branding staging

NEW

55

Introduces app branding migration, letting eligible orgs stage, preview, and deploy Identity Engine on an app-by-app basis while preserving Classic Engine customizations, with per-app rollback instead of org-wide rollback.

– Explains staging and rollback behavior but no UI path or API named

`changelog-20261020-2c40ee6d`

05 AI agent import from Snowflake, Workato, and Cursor

NEW

50

Okta adds support for importing and managing AI agents built in `Snowflake`, `Workato`, and `Cursor` directly through Okta, extending its AI agent identity governance to these platforms.

– Names three integrations but no mechanism or UI path given `changelog-20261020-2c40ee6d`

06 NFC badge authenticator for Adaptive MFA

NEW

50

Adds NFC authenticator support for orgs with Adaptive MFA, allowing frontline workers to tap an NFC badge and enter a PIN to sign in on shared Windows workstations.

– Describes the mechanism but no config key or flag given `changelog-20261020-2c40ee6d`

07 RADIUS Agent default client access change

BREAKING

50

RADIUS Agent 2.28 changes the default client access behavior — new installations no longer accept RADIUS connections without an explicitly configured client IP allowlist.

– States the breaking behavior but no migration steps given `changelog-20261020-2c40ee6d`

08 App-level FastPass sign-in button control

IMPROVED

45

Introduces user identification policy rules so admins can control whether the `Sign in with Okta FastPass` button appears on an app-by-app basis instead of a single org-wide setting.

– Names the exact UI element but not where to configure it `changelog-20261020-2c40ee6d`

09 OIN Community Integrations (Early Access)

NEW

40

Introduces OIN Community Integrations (Early Access), publishing integrations built from public APIs to the Okta Integration Network as Community Preview, advancing to Community Verified once proven in production.

– Describes lifecycle stages but no actionable steps given [changelog-20261020-2c40ee6d](#)

10 Passkey enrollment promotion nudge NEW 40

Adds a configurable passkey enrollment promotion nudge that prompts end users to enroll a passkey when signing in, with controls for re-prompt frequency and maximum skip count.

– Mentions configurable controls but not exact settings or keys

[changelog-20261020-2c40ee6d](#)

11 New provisioning integrations: Fleet End User Enrollment and Pipedrive

NEW 30

Adds provisioning support for the [Fleet End User Enrollment](#) and [Pipedrive](#) integrations.

– Bare name of two integrations with no further detail [changelog-20261020-2c40ee6d](#)

! BREAKING ON UPGRADE

! RADIUS Agent 2.28 changes the default client access behavior — new installations no longer accept RADIUS connections without an explicitly configured client IP allowlist.

WAS THIS USEFUL?



Twingate

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

RANK

Twingate is a identity-based access for users, services, and AI agents that deploys in minutes, scales to every resource, and finally lets you retire your VPN.

Twingate added HTTPS support to Privileged Access for Web Apps, removing browser certificate warnings and duplicate sign-ins for internal web apps.

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

01 HTTPS support in Privileged Access for Web Apps

IMPROVED

50

Twingate Privileged Access for Web Apps now supports HTTPS, eliminating browser warnings and second sign-ins for internal web apps. The feature covers every leg of the user path to internal web apps — including public and private CAs — without requiring per-app certificate management.

– Describes mechanism and CA scope but no exact config or UI path given.

changelog-20261020-982f698d

WAS THIS USEFUL?



Give feedback to the DNSFilter team so we can make more informed product decisions.
Powered by Canny.

DNSFilter rolled out multilingual support for block-page access requests and completed an automatic migration of branded admin subdomains to a new domain.

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

01 Branded subdomain migration to dnsfilter.io IMPROVED

55

Managed branded subdomains are automatically migrated from `app.dnsfilter.app` to `app.dnsfilter.io`, with no admin action required.

– Names exact old and new domains but migration is automatic, nothing to act on.

changelog-20261020-0d395246

02 Multilingual Request Access form on block pages NEW

48

The Request Access form on block pages now supports Spanish, French, and German, auto-detecting the user's browser language to display the appropriate translation.

– Describes behavior and languages but no config or admin control named.

changelog-20261020-0d395246

WAS THIS USEFUL?



Tailscale creates private WireGuard-based mesh networks with centralized identity and access controls.

This window's release focused entirely on Tailscale's Kubernetes operator, adding PeerRelay custom `UDP` ports and automatic re-authentication, per-tailnet Let's Encrypt account key sharing, broader image pull secret application, Helm resource name overrides, safer rolling egress updates, and a fix restoring in-cluster Ingress reachability.

←→ WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 Ingress traffic restored via `TS_SERVE_ALLOW_ALL_INTERFACES`

IMPROVED

80

Adds the `TS_SERVE_ALLOW_ALL_INTERFACES` environment variable, set by the operator on proxies for Ingress resources annotated with `tailscale.com/experimental-forward-cluster-traffic-via-ingress`, keeping those Ingress resources reachable from inside the cluster after Tailscale Serve on Linux was scoped to the Tailscale interface only.

– Names exact env var and annotation plus the regression it fixes

changelog-20261020-c5f6abb7

02 Helm resource name overrides for operator install NEW

65

Adds `nameOverride` and `fullnameOverride` Helm values to customize the names of resources the operator installs.

Customize Helm-installed resource names to avoid collisions in shared clusters.

```
yaml
nameOverride: my-tailscale
fullnameOverride: my-tailscale-operator
```

– Exact Helm keys with a runnable config example

changelog-20261020-c5f6abb7

THINNER COVERAGE BELOW

03 Custom UDP ports for PeerRelay resources NEW

55

Adds custom UDP port support for `PeerRelay` resources so peers that cannot reach the relay at its load balancer address have an alternative path.

– Explains purpose but no field or command for setting the port

04 Shared Let's Encrypt account key per tailnet NEW 50

Supports sharing a single persistent Let's Encrypt account key per tailnet to keep certificate renewals exempt from Let's Encrypt rate limits.

– States benefit but no config key or enabling step [changelog-20261020-c5f6abb7](#)

05 Rolling updates wait for egress Service routing IMPROVED 45

Rolling updates of egress proxies now wait for each new pod to route traffic to all of its egress Services before proceeding to the next pod.

– Clear before/after behavior but no flag to control it [changelog-20261020-c5f6abb7](#)

06 PeerRelay replicas re-authenticate after losing login IMPROVED 40

PeerRelay replicas re-authenticate after losing their login, such as when the device is deleted or its auth key expires.

– Names trigger conditions but not the recovery mechanism [changelog-20261020-c5f6abb7](#)

07 Image pull secrets applied to all ServiceAccounts IMPROVED 35

Applies image pull secrets to every ServiceAccount the operator creates, not only the operator ServiceAccount.

– Bare statement of scope change, no configuration surface named

[changelog-20261020-c5f6abb7](#)

08 PeerRelay resources use configured coordination server IMPROVED 35

PeerRelay resources now use the coordination server URL configured for the operator.

– Bare statement without further mechanism or config detail [changelog-20261020-c5f6abb7](#)

WAS THIS USEFUL?



iDenfy

Feedback



changelog · Oct 20, 2026

CHANGELOG ›

RANK

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

iDenfy's latest release adds MCP-based AI assistant access to accounts, launches a standalone Bank Card Verification flow inside a dedicated PCI-DSS v4.0.1 environment with Brazil CPF support in its iOS and Android SDKs, and introduces Persian (Farsi) RTL support, alongside smaller updates to KYC reporting, theming, filtering and KYB workflows.

←→ WHAT SHIPPED · 9 FEATURES most completely described first ? what's the number?

01 Standalone Bank Card Verification in PCI-DSS v4.0.1 environment

NEW

78

Runs Bank Card Verification as a standalone flow, independently of identity verification, inside a dedicated PCI-DSS v4.0.1 certified environment, and adds Bank Card Verification and Brazil CPF check support to the iOS and Android SDKs.

– Names flow, certification and the two SDKs it ships in `changelog-20261020-e3b99304`

02 MCP-based AI assistant access

NEW

65

Enables AI assistants, starting with Claude and Claude Code, to connect to iDenfy accounts over MCP, with a consent screen to control which data areas they can view or edit.

– Names the protocol and assistants but no setup steps `changelog-20261020-e3b99304`

03 New fields in KYC reports and results

NEW

62

Adds an Original language data section to KYC CSV reports, exporting document field values in their native script as separate columns, and adds a marital name field to KYC results for ID documents that carry a second surname.

– Names the exact report section and export mechanism Names the field but no usage context

`changelog-20261020-e3b99304`

THINNER COVERAGE BELOW

04 Persian (Farsi) RTL support in verification flow

NEW

52

Adds Persian (Farsi) language support with full RTL layout and automatic photo capture on the selfie step to the identity verification flow.

– Describes the feature but no config or toggle named `changelog-20261020-e3b99304`

05 **KYC theme defaults and button branding color** IMPROVED

50

Allows any KYC theme to be marked as the default for new sessions, and adds primary button branding color as a separate configurable option.

– Names two configurable options, no exact UI path `changelog-20261020-e3b99304`

06 **KYB AI report and ownership chain improvements** IMPROVED

45

KYB AI reports now look up companies in the registry step by step, and KYB ownership chains can now end at a company declared to have no further shareholders.

– Describes behaviour change but no interface detail `changelog-20261020-e3b99304`

07 **Continent-based country filtering in verifications list** NEW

40

Adds continent-based country filtering to the verifications list.

– Clear UI feature but no further detail given `changelog-20261020-e3b99304`

08 **KYB update requests shared by link** NEW

40

Enables KYB update requests to be shared by link.

– Clear capability but no mechanism or format specified `changelog-20261020-e3b99304`

09 **Brazil CPF checks collect date of birth** IMPROVED

30

Brazil CPF checks now collect a date of birth.

– Single sentence, no mechanism or scope given `changelog-20261020-e3b99304`

WAS THIS USEFUL?



RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

ConductorOne's biggest addition is C1AI expanding across findings remediation, request/review policy steps, and Slack, alongside a breaking change to how connector-synced resources are deleted via API, plus new secure secrets-sharing, custom campaign selectors, and local policies.

↳ **WHAT SHIPPED · 8 FEATURES** most completely described first [? what's the number?](#)

01 Connector-synced resources protected from API deletion

88

BREAKING

DELETE

`/api/v1/apps/{app_id}/resource_types/{resource_type_id}/resources/{id}` now returns an error when called on connector-synced resources instead of deleting them (previously deleted resources simply reappeared on the next sync); the API also now exposes a read-only `isManuallyManaged` field on app resources returned by the API. Automation relying on the old delete behavior must be updated.

Safely delete only manually managed app resources via API — connector-synced resources will now return an error instead of being silently removed.

```
$ curl -X DELETE
'https://<tenant>.conductor.one/api/v1/apps/<app_id>/resource_types/<resource_type_id>/resources/<id>' \
-H 'Authorization: Bearer <token>'
```

– Full endpoint path, field name, and runnable curl example [changelog-20261020-6d00b310](#)

02 Findings: AI remediation and triage actions

NEW

78

Application findings now include a **Remediate with AI** button that opens **C1AI** with the finding's type and target pre-filled, following a built-in remediation playbook for unowned service accounts, non-human identities, shadow apps, and connector sync failures, plus a new Findings tab on application detail pages listing open and in-progress findings with target, assignee, and filters for severity, type, and assignee; shadow apps are now triageable findings supporting assign, snooze, suppress, accept risk, reopen, and **Remediate with AI** actions from the Shadow apps page.

:

– Names button, playbook scope, and new findings UI Names exact triage actions and page, no command

changelog-20261020-6d00b310

03 **Policy enhancements: C1AI steps and scoped local policies**

NEW

70

C1AI can now be added as an assignable workflow step in request and review policies, letting the agent evaluate each request using C1's built-in guidance plus custom instructions entered per step, and local policies can now be created directly from the policy menu on an app or entitlement, remaining bound to that single app or entitlement outside the global policy library.

Assign C1AI as an automated evaluator step in a review policy so the agent screens every request that reaches it before a human approver is involved.

 In the C1 console, go to Policies > select or create a review policy > Add step > Assignee: C1AI > optionally enter custom instructions > Save.

– Exact console navigation path given as example Clear UI starting point, limited further mechanism detail

changelog-20261020-6d00b310

04 **Campaign entitlement management and scoping**

NEW

65

Campaign templates can now use batch **Add template entitlements** and **Remove template entitlements** API calls, each accepting up to 200 entitlements, to build selections incrementally without overwriting earlier choices, and access review campaigns can be scoped using custom selectors that filter by app, resource type, risk level, and compliance framework, automatically pulling in newly matching entitlements when a campaign is prepared.

– Named API calls and limit but no endpoint path given Names filter dimensions but no exact UI path

changelog-20261020-6d00b310

05 **Access request links now require explicit submission**

BREAKING

61

Access request links, including older auto-submit links generated from Copilot or Jira, no longer submit requests automatically — they only prefill the form. Jira templates, runbooks, or automations expecting a link to create a request directly must be updated to expect an explicit user submit.

– Clear before/after and migration note, no exact field names

changelog-20261020-6d00b310

06 Secure secrets-sharing from C1 NEW

58

Adds secure secrets-sharing directly from ConductorOne: content is encrypted in-browser before upload, recipients authenticate via SSO or a one-time magic link, and administrators can view metadata and audit logs for all secrets shared across the tenant.

– Describes mechanism but no exact UI path or command [changelog-20261020-6d00b310](#)

07 Character limits on Text and Text area form fields NEW

50

Adds an optional character-limit setting on **Text** and **Text area** fields in request forms, preventing submission of values that exceed the limit.

– Names fields and behavior but no exact config path [changelog-20261020-6d00b310](#)

08 C1AI assistant natively in Slack NEW

41

C1AI is now available natively in Slack with rich interactive responses — buttons, forms, and pickers — letting users complete tasks without leaving the conversation.

– No setup steps or commands given beyond description [changelog-20261020-6d00b310](#)

↳ ALSO FROM THESE RELEASES

.

↳ BREAKING ON UPGRADE

- ! **DELETE** `/api/v1/apps/{app_id}/resource_types/{resource_type_id}/resources/{id}` now returns an error when called on connector-synced resources; previously it deleted them and they reappeared on the next sync. Any automation that depended on this deletion succeeding must be updated.
- ! Access request links (including older auto-submit links from Copilot or Jira) no longer submit requests automatically — they now only prefill the form. Jira templates, runbooks, or automations that expected a link to create a request must be updated to expect an explicit user submit.

WAS THIS USEFUL?



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

Auth0 extended Universal Login's phone and identifier experience with new country controls and UI, and introduced a guided, visual builder for Tenant Access Control (ACL) policies.

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

01 Guided Tenant ACL policy builder and management NEW

80

Adds a Guided Rule Builder under **Security > Access Control** for constructing complex Tenant ACL policies via step-by-step prompts for scope, entity, and action selection. Pairs it with Visual Rule Sequencing, an interactive canvas to reorder ACL rule priorities and inspect exact execution order, and Instant Status Toggles to enable or disable individual Tenant ACL rules with a single click from the dashboard.

Build and sequence a Tenant ACL rule without hand-writing policy syntax, then toggle it live immediately.

📌 In the Auth0 dashboard, navigate to Security > Access Control, use the Guided Rule Builder to set scope, entity, and action, drag the new rule to the correct position in the Visual Rule Sequencing canvas, then use the Instant Status Toggle to enable it.

– Exact dashboard path plus named sub-features and steps. `changelog-20261020-9bca9ebb`

02 Phone country allow/deny lists for Universal Login NEW

60

Adds tenant-level include/exclude lists, configured via tenant settings, to control which phone countries are supported during phone entry in Universal Login.

Restrict Universal Login phone entry to a specific set of allowed countries at the tenant level.

📌 In the Auth0 dashboard, go to tenant settings, open the phone country rules section, and add the countries you want to the include list.

– Names the config surface but not an exact key or API. `changelog-20261020-9bca9ebb`

THINNER COVERAGE BELOW

03 Updated phone entry UI in Universal Login NEW

55

Adds an updated phone entry UI with country selection, regional number formatting, a numeric keyboard, and phone number suggestions on supported devices. Available in Early Access on all plans.

– Concrete UI changes listed but no exact setup path given. changelog-20261020-9bca9ebb

04 **Theme API controls for identifier presentation** NEW

40

Adds theme API controls for identifier presentation in Universal Login, letting type-specific **Continue with** buttons be shown instead of a single unified identifier field.

– Describes the toggle but gives no steps or exact key. changelog-20261020-9bca9ebb

ALSO FROM THESE RELEASES

- .
- .
- .

WAS THIS USEFUL?



Nudge Security

Feedback



changelog · Oct 20, 2026

CHANGELOG

RANK

Nudge Security is a SaaS security platform that discovers shadow IT assets, manages identity governance, and hardens security posture.

Nudge Security added custom SSO domain tracking that lets the browser extension detect SSO logins without an IdP connection, powering new inventory filtering and auto-approval.

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

01 Custom SSO domain tracking and auto-approval NEW 80

Custom SSO domains can now be configured under **Settings > Browser extension**, letting the browser extension recognise logins through those domains and auto-fill the SSO provider in the app inventory without requiring an IdP connected-app. This also enables filtering the app inventory to show which apps use SSO based on detected logins through those domains, and enables auto-approval of apps that users sign into via the configured custom SSO domains.

Surface SSO coverage across your app inventory without requiring a fully configured IdP connected-app for every tool.

In the Nudge Security console, go to Settings > Browser extension and add your custom SSO domains (e.g. **ss0.yourcompany.com**). The browser extension will then detect logins through those domains and automatically fill in the SSO provider for each matching app in your inventory.

– Names exact config path and three behaviors, with setup steps

changelog-20261020-6777db71

WAS THIS USEFUL?



Ping Identity

[Feedback](#)[changelog · Oct 20, 2026](#)[CHANGELOG ↗](#)

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 Advanced Identity Cloud release adds configurable [Entra ID/SMTP](#) email providers with test-send support, customizable hosted-page navigation menus, and an option to skip XML signature verification on [SAML](#) metadata import, alongside smaller admin console and journey designer usability improvements.

WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 Entra ID and custom SMTP email provider configuration NEW

75

Adds ability to configure Microsoft Entra ID or custom SMTP email providers and send test emails from the Advanced Identity Cloud admin console. Reach it via Email Provider settings, select [Microsoft Entra ID](#) or [Custom SMTP](#), fill in the connection details, then click [Send Test Email](#) to confirm delivery.

Route outbound identity emails through Microsoft Entra ID or a custom SMTP server and verify delivery with a test email before going live.

📌 In the Advanced Identity Cloud admin console, go to Email Provider settings, select [Microsoft Entra ID](#) or [Custom SMTP](#), fill in the connection details, then click [Send Test Email](#) to confirm.

– Exact admin console navigation and test-email steps given. [changelog-20261020-b3fc67d8](#)

02 Navigation profiles for hosted account pages NEW

75

Adds a customizable hosted account pages navigation menu via navigation profiles, enabling different menus for different groups of users. Configure it in the admin console under Hosted Pages > Navigation Profiles, create a new profile, assign it to a target user group, and set the menu items to expose.

Tailor the hosted account portal menu so that, for example, admin-tier users see account management links that standard users do not.

📌 In the admin console, go to Hosted Pages > Navigation Profiles, create a new profile, assign it to the target user group, and configure the menu items to expose.

– Exact admin console path to create and assign profiles. [changelog-20261020-b3fc67d8](#)

03 Journey canvas multi-select and pan controls IMPROVED

65

Improves journey canvas interactions by adding the ability to select multiple journey nodes by holding Shift and clicking each node, and to pan the canvas by holding the space bar and dragging with the left mouse button.

– Exact key combinations given for both gestures. [changelog-20261020-b3fc67d8](#)

THINNER COVERAGE BELOW

04 Telephone number user property format NEW

55

Adds `telephone number` as a user property format with optional server-side validation, providing separate country code and phone number fields in hosted pages and the admin console.

– Describes fields and validation but no exact config path. [changelog-20261020-b3fc67d8](#)

05 Next-generation scripting engine made default IMPROVED

40

Makes the next-generation scripting engine the default when creating scripts in the Advanced Identity Cloud admin console.

– States the change but not what differs from prior engine. [changelog-20261020-b3fc67d8](#)

06 Skip XML signature verification on SAML metadata import NEW

38

Adds support for skipping XML signature verification when importing remote SAML 2.0 provider metadata.

– No UI path or flag given for enabling this option. [changelog-20261020-b3fc67d8](#)

07 Multi-value custom node properties NEW

38

Adds multi-value selection from predefined options when configuring custom node properties in the node designer.

– No navigation path or example given. [changelog-20261020-b3fc67d8](#)

08 Search and sort for AI agent identities NEW

28

Adds search and sort capabilities for the list of AI agent identities.

– Bare description, no UI path or filter details given. [changelog-20261020-b3fc67d8](#)

WAS THIS USEFUL?



All the latest releases from the 1Password team

1Password's SaaS Manager added CSV exports for AI consumption data, support for multiple overlapping budgets per application, GitHub Copilot usage tracking, and SCIM-based account provisioning for OpenAI Platform, alongside several smaller refinements to dashboards and workflows.

WHAT SHIPPED • 8 FEATURES most completely described first ? what's the number?

01 CSV export for AI consumption data NEW 81

SaaS Manager now supports exporting AI consumption data as CSV from the Spend Explorer's AI Consumption tab and from the application-level AI consumption tab. The export respects active filters and the selected date range, with an optional Group by month toggle.

Pull export-ready AI consumption totals for a cost review — filter by date range in the spend explorer, then download the CSV grouped by month.

Steps: In SaaS Manager, open Spend Explorer > AI Consumption tab. Set the desired date range and filters. Click Export and select CSV. To group by month, enable the Group by month option before exporting.

– Exact UI path and export steps given, with a runnable usage example

1Password SaaS Manager #2075 (2026-10-1...

02 Multi-budget support per application NEW 61

SaaS Manager now supports creating multiple budgets for the same application with overlapping date ranges, each tracking spend, threshold alerts, and runway calculations independently.

– Describes behavior and scope but no concrete UI steps or limits

1Password SaaS Manager #2075 (2026-10-1...

03 New ID merge fields for Workflow API request steps IMPROVED

58

API request steps in New Workflows now support additional ID merge fields: `person`, `application`, `task`, `user account`, and related person and application IDs.

– Names the specific merge fields but no example workflow shown

1Password SaaS Manager #2075 (2026-10-1...

04 SCIM provisioning for OpenAI Platform NEW

54

SaaS Manager adds SCIM account management for OpenAI Platform organizations on the Enterprise plan, supporting automated account provisioning and deprovisioning.

– Names protocol and plan tier but no endpoint or setup detail

1Password SaaS Manager #2075 (2026-10-1...

05 AI Consumption dashboard refinements IMPROVED

53

Several small refinements landed in the AI Consumption dashboards: spike details now fall back to token metrics when spend data is unavailable, covering sources such as API keys; the `License Not In Use` widget now focuses on unengaged accounts with confirmed paid licenses to cut noise from free or unknown-license accounts; and AI Consumption integration descriptions now explain which dashboard metrics each provider supports.

– Names all three constituent changes but each is thinly described

1Password SaaS Manager #2075 (2026-10-1...

06 GitHub Copilot AI consumption tracking NEW

52

AI Consumption now adds support for GitHub Copilot, enabling import of user-level credits and token usage as well as application-level credits broken down by model.

– Names the integration and data granularity but no setup steps

1Password SaaS Manager #2075 (2026-10-1...

07 Provider project/workspace ID mapping in Projects IMPROVED

50

SaaS Manager Projects mapping now exposes provider project and workspace IDs and names, using stable provider IDs that preserve usage history across renames.

– Explains mechanism but no concrete navigation or example

1Password SaaS Manager #2075 (2026-10-1...

08 Rippling personal email in person profiles IMPROVED

43

Rippling personal email addresses now appear in SaaS Manager person profiles and are available for use in onboarding workflows.

– Simple, clearly described addition with no further mechanism

WAS THIS USEFUL?



DigiCert CertCentral

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

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

CertCentral adds a **DigiCert default (recommended)** chain menu option, cleans Apache certificate bundles, and deprecates the **disable_ct** API parameter.

- › Adds **DigiCert default (recommended)** option to the **Default intermediate chain** product menu so operators can distinguish a DigiCert-chosen chain from one they explicitly selected; the **ca_cert_id** API parameter continues to override this setting per-request.
- › Apache download (**DigiCertCA.crt**) for G5 and G2 TLS certificates will no longer include the unneeded self-signed root — only the cross-signed root remains, removing the need for manual post-download cleanup scripts.
- › The **disable_ct** parameter on certificate order requests (including **Order Basic OV**) is deprecated; CertCentral now ignores it and always logs issued TLS certificates to CT logs regardless of the value passed.

↳ BREAKING ON UPGRADE

- ! Scripts that remove a certificate by position in the Apache **DigiCertCA.crt** file, or that expect a specific number of certificates in the file, will break after October 20, 2026 — update scripts to remove the self-signed root only if present and always retain the cross-signed root.
- ! The **disable_ct** parameter is ignored as of this change: requests that set **disable_ct** to **true** to suppress CT logging will no longer have that effect — certificates will be logged to CT logs unconditionally.
- ! Starting November 12, 2026, new browser-trusted TLS orders, renewals, reissues, and duplicates will default to G5 root hierarchies; accounts relying on Global G2 or G3 intermediate chains must explicitly select those chains via the **Default intermediate chain** menu or the **ca_cert_id** API parameter before that date.

COVERAGE

NEW supported platform [G5 TLS root hierarchy transition](#)

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

WAS THIS USEFUL?



Veeam Kasten

[Feedback](#)

9.0.8

[NOTES](#)[RANK](#)

Veeam Kasten is a Kubernetes-native platform for backing up, restoring, and migrating containerized applications and their persistent data.

Kasten K10 9.0.8 adds restartless per-service log-level control, automatic image-pull-secret propagation for private registries, configurable per-policy action timeouts, KubeVirt filesystem-disk backup, and SBOM/OpenVEX publishing for vulnerability-scanner suppression, alongside several breaking platform and validation changes.

←→ WHAT SHIPPED · 11 FEATURES most completely described first ? what's the number?

01 SBOM and OpenVEX bundles for vulnerability suppression NEW

93

Veeam Kasten now publishes an SBOM bundle (`sboms-9.0.8.tar.gz` , SPDX format) covering all images and an OpenVEX document per release (`vex-9.0.8.openvex.json`); passing `--vex vex-9.0.8.openvex.json` to `grype` (v0.68.0+) or `trivy image` automatically suppresses VEX-marked false-positive findings for Go modules whose vulnerable code isn't present in Kasten binaries, with `--show-suppressed` (table output only) listing suppressed results alongside the rest. Published SBOMs identify each image's root package by image name, Kasten version, image digest, and package URLs (`pkg:oci` and `pkg:generic`) for accurate VEX matching; older grype versions can apply suppression via config `ignore: - vex-status: not_affected / fixed` .

Scan a Kasten image SBOM with grype and suppress VEX-confirmed false positives in one step.

```
$ tar -xzf sboms-9.0.8.tar.gz
grype sbom:dashboardbff_9.0.8.spdx --vex vex-9.0.8.openvex.json
```

Scan a Kasten container image directly with Trivy and suppress VEX-marked findings.

```
$ trivy image --vex vex-9.0.8.openvex.json gcr.io/kasten-
images/dashboardbff:9.0.8
```

Apply VEX suppression in grype versions older than 0.68.0 via the grype config ignore rules.

yaml

```
ignore:
  - vex-status: not_affected
  - vex-status: fixed
```

– Exact filenames, flags and config snippet with runnable commands 9.0.8

02 Restartless per-service log-level control

NEW

91

The `logLevelOverrides` Helm map sets per-service log level (`info` or `debug`), overriding the global `logLevel` without restarting pods, gated by `features.restartlessLogLevelChange` (now enabled by default on new installs); remove an override with `--set-string logLevelOverrides.<service>=`. Services read the `loglevel` key of the `k10-config` ConfigMap at runtime so changes take effect within seconds, and new operation pods inherit the log level of the service that creates them (explicit `LOG_LEVEL` settings and CLI options still take precedence). On upgrade with `--reuse-values` the feature must be enabled explicitly — doing so restarts all service pods — and once enabled Kanister follows the same levels as other services, treating any `logLevel` other than `info` / `debug` as `info`.

Enable restartless log-level changes and set the catalog service to debug without bouncing all pods — useful for live troubleshooting in production.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io --reuse-values \
  --set features.restartlessLogLevelChange=true \
  --set-string logLevelOverrides.catalog=debug
```

Tune per-service log verbosity at runtime without restarting Kasten — useful when debugging a specific service such as the catalog in production.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io --reuse-values \
  --set features.restartlessLogLevelChange=true \
  --set logLevelOverrides.catalog=debug
```

– Exact Helm keys, ConfigMap behaviour and upgrade caveats all given 9.0.8

03 Configurable action-phase timeouts, global and per-policy

NEW

84

Per-operation maximum timeout Helm keys (in minutes) now cover Blueprint backup, Generic Storage Backup per PVC, volume snapshot export per PVC, Blueprint/storage

restore per PVC, Blueprint delete, and exported volume data retirement, and each can be overridden per policy via annotations `k10.kasten.io/backup-timeout`, `k10.kasten.io/restore-timeout`, `k10.kasten.io/retire-timeout`, and `k10.kasten.io/blueprint-resource-hooks-timeout` set on the policy object.

Override the default backup and blueprint hook timeouts on a long-running database policy without changing the global K10 defaults.

```
$ kubectl annotate policy my-db-backup-policy \
  k10.kasten.io/backup-timeout=4h \
  k10.kasten.io/blueprint-resource-hooks-timeout=30m \
  -n kasten-io
```

Override the backup and restore timeouts for a latency-sensitive policy without touching the global defaults.

```
$ kubectl annotate policy my-backup-policy \
  k10.kasten.io/backup-timeout=2h \
  k10.kasten.io/restore-timeout=3h \
  --namespace kasten-io
```

– All six Helm timeout areas and four annotation names given with a runnable example

9.0.8

04 Image pull secret propagation to app namespaces NEW

75

The `global.copyImagePullSecretsToAppNamespaces` Helm key, paired with `global.imagePullSecrets`, automatically copies Veeam Kasten's image pull secrets into application namespaces during Blueprint actions and Generic Storage Backup, with automatic cleanup once the operation finishes; required for pulling from private registries on Kubernetes 1.35 and later.

Propagate image pull secrets to app namespaces for clusters running Kubernetes 1.35+, where K10 must access private registries during Blueprint and Generic Storage Backup.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io --reuse-values \
  --set global.copyImagePullSecretsToAppNamespaces=true
```

Enable automatic temporary image pull secret propagation to app namespaces during offline backup/restore, so pods can pull sidecar images without manual secret management.

```
$ helm upgrade k10 kasten/k10 --set
global.copyImagePullSecretsToAppNamespaces=true --set
global.imagePullSecrets=<your-secret-name>
```

– Named Helm keys with exact use case and runnable command 9.0.8

05 IronBank image management in k10tools NEW

75

`k10tools` gains `ironbank image list` and `ironbank image copy --dst-registry` subcommands for listing and mirroring IronBank image variants into a private registry, for managing air-gapped deployments.

Mirror IronBank-variant Kasten images to a private registry in a disconnected environment.

```
$ docker run --rm -v ~/.docker:/home/kio/.docker gcr.io/kasten-
images/k10tools:9.0.8 ironbank image copy --dst-registry
repo.example.com
```

– Exact subcommands and a runnable mirroring command product docs

06 ActionPodSpecBinding schema validation and binding precedence

70

BREAKING

`ActionPodSpecBinding` objects are now validated against a schema on create/update: unknown fields under `spec` are rejected and not stored, and resource quantities must be an integer or quoted string — a bare decimal like `cpu: 0.5` is rejected in favor of `cpu: '0.5'` or `cpu: 500m` (existing objects keep working until next updated). The `repository-server` datamover pod and `delete-collection` pod now prefer an `ActionPodSpecBinding` in the protected application's namespace over one in the Veeam Kasten install namespace, and binding in the install namespace for these pods is deprecated and will be removed in a future release.

– Exact field/format rules given but no remediation command 9.0.8

07 Repository operations concurrency limiter NEW

66

The `limiter.repositoryOperationsPerCluster` Helm value (default `3`) caps the number of concurrent background repository maintenance runs and storage usage scans per cluster, reducing storage I/O contention during busy backup windows.

Cap concurrent background repository maintenance runs cluster-wide to avoid storage I/O saturation during busy backup windows.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io --reuse-values \
  --set limiter.repositoryOperationsPerCluster=3
```

Limit how many repository maintenance and storage-scan jobs run in parallel to avoid overwhelming object-store bandwidth.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io --reuse-values \
  --set limiter.repositoryOperationsPerCluster=2
```

Limit concurrent background repository maintenance to 1 run on a cluster with constrained storage I/O bandwidth.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io --reuse-values \
  --set limiter.repositoryOperationsPerCluster=1
```

– Named key and default plus multiple runnable examples 9.0.8

THINNER COVERAGE BELOW

08 Platform requirement changes: Kubernetes and CPU

BREAKING

55

Veeam Kasten now supports Kubernetes 1.36, while support for Kubernetes 1.32 has been removed (OpenShift 4.19 clusters remain supported); on x86-64 clusters it now requires an x86-64-v3 CPU (Intel Haswell / AMD Excavator) or newer, dropping support for older x86-64 CPUs.

– Clear before/after versions but no remediation steps 9.0.8

09 Ephemeral PVC size margin for staged ImageStream images **NEW**

45

A new ephemeral PVC size margin Helm key (a decimal fraction, e.g. **0.1** for 10%, capped at **10**) controls the size of PVCs staged for ImageStream images, computed as **PVC size = (image bytes) * (1 + margin) + 64 MiB**.

– Formula given but exact key name and command missing product docs

10 KubeVirt filesystem VM disk backup **NEW**

40

Veeam Kasten can now back up Filesystem VolumeMode KubeVirt VM disks using the `VirtualMachineBackup` API.

– Names the API but gives no usage detail or command `9.0.8`

11 RBAC csinodes permission for self-managed deployments

35

BREAKING

Deployments using self-managed service account RBAC (`rbac.create=false`) must now grant `list` permission on `csinodes` in the `storage.k8s.io` API group to the appropriate service account.

– States the exact permission needed but no manifest or command `9.0.8`

➡ BREAKING ON UPGRADE

- ! When upgrading an existing installation with `--reuse-values`, `features.restartlessLogLevelChange` must be explicitly enabled; enabling it restarts all service pods.
- ! When using `--reuse-values` with `logLevelOverrides`, supplying `null` for an override causes Helm to restore the previous value rather than removing it; use `--set-string logLevelOverrides.<service>=` (empty string) to actually remove an override.
- ! An unsupported log level value, or a service name that cannot read a level from the ConfigMap, causes the install or upgrade to fail.
- ! Support for Kubernetes 1.32 is removed (OpenShift 4.19 clusters remain supported).
- ! On x86-64 clusters, Veeam Kasten now requires x86-64-v3 CPU (Intel Haswell / AMD Excavator) or newer; older x86-64 CPUs are no longer supported.
- ! Deployments using self-managed service account RBAC (`rbac.create=false`) must grant `list` permission on `csinodes` in the `storage.k8s.io` API group to the appropriate service account.
- ! `ActionPodSpecBinding` objects are now validated against a schema; unknown fields under `spec` are rejected and not stored. An existing object with an invalid resource quantity is not corrected by the upgrade.
- ! Resource quantities in an `ActionPodSpecBinding` must be an integer or a string — a bare decimal (e.g. `cpu: 0.5`) is rejected; use `cpu: '0.5'` or `cpu: 500m` instead. Existing objects with bare decimals continue working until next updated.
- ! Using an `ActionPodSpecBinding` in the Veeam Kasten install namespace to configure the `repository-server` datamover Pod or the `delete-collection` Pod is deprecated and will be removed in a future release; bind it in the protected application's namespace instead.
- ! With `features.restartlessLogLevelChange` enabled, Kanister follows the same log levels as other Veeam Kasten services; a `logLevel` other than `info` or `debug` (e.g. `error`) causes Kanister to log more after the upgrade. An upgrade with `--reuse-values` keeps prior defaults, so set `features.restartlessLogLevelChange=true` explicitly if desired.
- ! `ActionPodSpec` objects are now validated on create/update; unknown fields under `spec` are no longer stored and `kubect1` rejects them. A bare decimal CPU value such as `cpu:`

`0.5` is rejected — write it as `cpu: "0.5"` or `cpu: 500m`.

- ! The `repository-server` datamover pod and the `delete-collection` pod now prefer an `ActionPodSpecBinding` in the protected application's namespace over one in the Veeam Kasten install namespace, which may change behaviour for existing deployments using install-namespace bindings.
- ! Kasten deployments using self-managed RBAC (`rbac.create=false`) must grant `list` permission on `csinodes` in the `storage.k8s.io` API group to the appropriate service account.
- ! On x86-64 clusters, Veeam Kasten now requires an x86-64-v3 CPU (Intel Haswell / AMD Excavator) or newer; older x86-64 CPUs are no longer supported.
- ! Support for Kubernetes 1.32 has been removed (OpenShift 4.19 clusters continue to be supported).
- ! Enabling `features.restartlessLogLevelChange` changes Kanister log-level behaviour: a `logLevel` other than `info` or `debug` is treated as `info`, so setting `logLevel` to `error` will cause Kanister to log more after upgrade. Upgrades using `--reuse-values` do not automatically inherit the new default; set `features.restartlessLogLevelChange=true` explicitly.

WAS THIS USEFUL?



GOVERNANCE, RISK & COMPLIANCE

RegScale

Feedback



CLI 6.54.0 (2026-10-19)

BLOG ↗

RANK

RegScale is a compliance automation platform that helps organizations manage security controls, risk, and regulatory frameworks like NIST, FedRAMP, and SOC 2 through continuous monitoring.

RegScale's CLI 6.54.0 release adds a background prefetch option to speed up slow vendor scans and expands JSON run summaries with performance metrics.

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

01 Background prefetch for vendor scanner data NEW 72

A new `scannerPrefetch` setting fetches vendor findings and assets in the background while earlier ones are still processed, speeding up scans against slow vendor APIs.

Enable it with `scannerPrefetch: true` in config.

Enable background prefetching to speed up scans when your vendor APIs respond slowly.

```
yaml
scannerPrefetch: true
```

– Names exact config key and purpose, with usage example. CLI 6.54.0 (2026-10-19)

THINNER COVERAGE BELOW

02 Run duration and peak memory in JSON output IMPROVED 48

The `--output json` run summary now reports run duration and CLI peak memory use.

– Names flag and fields but no example shown. CLI 6.54.0 (2026-10-19)

WAS THIS USEFUL?



EVALUATE

Perplexity numbat

[Feedback](#)

v0.3.0

[NOTES ↗](#)

Visibility into AI agent activity on endpoints, with on-device detection, optional pre-action blocking, and forensic reconstruction.

Adds `rules test --json` for automation, retained tool I/O for CEL rules and exports, opt-in POSIX compound-command enforcement, and improved subagent attribution.

↳ [GET THIS VERSION](#)

```
$ git clone --branch v0.3.0 https://github.com/perplexityai/numbat.git
# already have the repo? check out this version:
$ git checkout v0.3.0
```

- › Adds `rules test --json` flag to `numbat rules test` for machine-readable rule-test output suitable for automation pipelines.
- › Adds `--content full --content-scope messages` flags to scope full-content exports to message bodies only, now that `--content full` also captures redacted tool arguments and results.
- › Retains tool inputs and results in the event record so CEL rules and content exports can act on them, including MCP tool inputs and browser context.
- › Adds opt-in enforcement for statically resolvable POSIX compound commands via rules marked `enforce: true`, extending blocking coverage beyond single-command pre-action hooks.
- › Improves subagent attribution in timeline and event records, preserving subagent identity and correlation across capture paths.

↳ BREAKING ON UPGRADE

- ! Output schema is now `0.5.0`; update any downstream validators that enforce an exact version or closed schema.
- ! `--content full` now includes redacted tool arguments and results (source code, file contents, patches); use `--content full --content-scope messages` to restore message-body-only export behaviour.
- ! Rerun `numbat hook install` for each installed agent with the same options to update generated hooks and plugins after upgrading.

COVERAGE

[NEW](#) integration Muse Code live-hook and at-rest session support

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

WAS THIS USEFUL?



GOVERN

ToolHive

Feedback



v0.55.0

NOTES ↗

RANK

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

ToolHive v0.55.0 strengthens MCP server security and connectivity with TLS/SPIFFE listener support, group-scoped virtual MCP upstream filtering, Pi LLM gateway integration, and a new Claude Desktop connector.

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

01 **Claude Desktop client connector** NEW 70

Adds `thv client connector` for Claude Desktop, enabling automatic client wiring.

– Runnable subcommand named, behaviour briefly described v0.55.0

02 **Pi support in LLM gateway** NEW 65

Adds support for Pi in the LLM gateway setup and honors the `PI_CODING_AGENT_DIR` environment variable for Pi MCP and skills configuration.

– Names env var but not full config details v0.55.0

03 **DCR client re-registration via rotation key** NEW 65

Adds `dcrConfig.rotationKey` to re-register upstream DCR clients that have been evicted.

– Config key named but mechanism of eviction unexplained v0.55.0

04 **Group-based upstream filtering for vMCP** NEW 60

Exposes a group-based upstream filter via the vMCP CRD, letting operators scope which upstreams a virtual MCP session can reach.

– Names CRD mechanism and effect but no field names v0.55.0

THINNER COVERAGE BELOW

05 **RBAC restriction on pods/attach** IMPROVED 55

Restricts `pods/attach` RBAC permission to stdio transport only, tightening the least-privilege posture.

– Names permission and scope but no migration detail v0.55.0

06 **TLS and SPIFFE listener configuration** NEW 45

Adds TLS listener config and SPIFFE X.509 authority capability for securing MCP server endpoints.

– States capability but no config keys or flags given v0.55.0

07 Scopeless upstream token refresh IMPROVED 35

Allows omitting scope on upstream token refresh, simplifying token exchange flows.

– Brief description, no endpoint or flag named v0.55.0

WAS THIS USEFUL?



Backslash detects and prioritizes reachable vulnerabilities in cloud-native application code and dependencies.

Backslash introduced Local Model Discovery, extending its AI inventory and governance controls to shadow AI models running locally on endpoints.

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

01 Local Model Discovery for shadow AI NEW 83

Backslash now detects AI models installed on macOS and Windows endpoints via Ollama, LM Studio, Jan, GPT4All, the Hugging Face cache, MLX, LocalAI, and llama.cpp, with no developer setup required. Discovered local models are surfaced in the existing AI inventory alongside remote models so the same allow and deny list controls apply to endpoint-resident models; the system reports a permission gap when a model directory cannot be read (distinguishing that from an inspected directory with no models), deduplicates multiple paths pointing to the same model into a single inventory record, filters out files that merely resemble models, and ties each discovered model to its specific endpoint even when no AI agent is attached.

Review all local models discovered across enterprise endpoints and identify any policy violations before applying allow/deny controls.

In Backslash, go to AI Inventory > filter by local designation > review each model's risk details and policy violations > add to allow or deny list.

- Names eight runtimes and behaviour, but no API/flag, only UI path.

Local Model Discovery: Govern Shadow AI...

WAS THIS USEFUL?



Cloudflare

Feedback



changelog · Oct 20, 2026

CHANGELOG ↗

RANK

Cloudflare is a connectivity cloud platform delivering internet security, performance, and developer infrastructure services.

Cloudflare shipped two new Workers AI models (Clairic extraction and DeepSeek v4.1 Flash), step-up authorization for the API MCP server, expanded Precursor bot-detection modes with interstitial controls, Cron Triggers support in Local Explorer, a new `cf docs search` CLI command, and the ability to buy additional tunnels from the dashboard.

↳ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Workers AI models Clairic and DeepSeek v4.1 Flash NEW

78

Adds two new Workers AI models: `@cf/cloudflare/clairic`, an extraction model priced at \$0.10/M input tokens, and `@cf/deepseek-ai/deepseek-v4.1-flash-1017`, a 552B parameter multimodal mixture-of-experts model.

Extract typed structured values from unstructured invoice text using Clairic — avoiding brittle regex or LLM prompt engineering for fields like totals, dates, and contact lists.

```
$ curl
https://api.cloudflare.com/client/v4/accounts/<account_id>/ai/run
/@cf/cloudflare/clairic \
-H 'Authorization: Bearer <token>' \
-H 'Content-Type: application/json' \
-d '{
  "model": "clairic",
  "state": "Invoice INV-2291 from Northwind Supply, issued
2026-09-02, net 30. 4 pallets at $250.00 each, shipping $80.00.
Contact: Dana Ruiz, Priya Shah.",
  "fields": {
    "invoice_number": {"type": "text"},
    "issued": {"type": "date"},
    "due_date": {"type": "date", "description": "Issue date
plus the net payment days"},
    "total": {"type": "number", "decimals": 2, "unit": "USD",
"description": "Amount due including shipping"},
    "contacts": {"type": "list", "items": "text"}
  }
}'
```

– Named models with pricing and a runnable API example

changelog-20261020-55200c0d

02 Step-up authorization in API MCP server

NEW

75

The Cloudflare API MCP server now supports step-up authorization: clients can connect with minimal scopes and approve additional scopes on demand. When a tool call needs a scope not yet granted, the server returns an `insufficient_scope` challenge naming the required scope; choosing `Minimal` on the consent page starts the session with least privilege.

– Mechanism explained but only a consent-page UI path, no command

changelog-20261020-55200c0d

THINNER COVERAGE BELOW

03 cf docs search command

NEW

55

The Cloudflare CLI adds a `cf docs search` command to find documentation from the terminal, e.g. `cf docs search "rate limit requests with WAF"`, useful for quick lookups during incident response or agent-assisted scripting.

Find documentation on rate limiting from the terminal without opening a browser — useful for quick lookups during incident response or agent-assisted scripting.

```
$ cf docs search "rate limit requests with WAF"
```

– Runnable command but minimal description of scope [changelog-20261020-55200c0d](#)

04 Buy additional tunnels from dashboard NEW

35

Additional tunnels can now be purchased from the Tunnels page in the Cloudflare dashboard.

– Only a UI path with no further mechanism given [changelog-20261020-55200c0d](#)

WAS THIS USEFUL?



// END OF TRANSMISSION ■

The Cyber Toolchain · curated daily · October 20, 2026

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

► **Subscribe**

Feedback on this issue