THE CYBER TOOLCHAIN

NO. 125

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 30, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 159 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 57 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's kill switch cuts one user's MCP client session dead without knocking over every other agent on the ıt. GitLab's Orbit graph goes GA, so you can ask which services a

vim ? help

vulnerability actually reaches instead of grepping. Snyk now runs its header, cookie and TLS checks against every Extra Host, not just the primary URL.

RESPOND

Cut off one compromised agent session without halting the rest of your automation

Aembit revokes a named user's MCP client session on demand, targeted or in bulk and callable from the API, so a suspected agent compromise doesn't force a tenant-wide credential rotation. Runlayer 2.0.0 closes cross-session manipulation and IP-spoofing gaps in the same layer, which is what makes per-session revocation meaningful rather than theatre.

Aembit · Runlayer

IDENTIFY

Ask which running services a vulnerability actually touches

Orbit is now GA for Premium and Ultimate: a live graph linking code, pipelines, vulnerabilities and ownership, queryable in plain language or by exact regex, with MCP wired to the CLI. Blast-radius questions that previously meant tracing imports by hand and chasing down a CODEOWNERS file become a single query.

GitLab

PROTECT

Scan the hosts your app actually serves, not just the one you typed in

Snyk API & Web applies security header, cookie attribute and SSL/TLS checks to every Extra Host on a target — the vhosts and alternate domains that quietly diverge from the primary config. Prowler's AWS security group check now catches rules that only become open exposure when their CIDRs are combined, a class of misconfiguration that per-rule evaluation reads as clean.

Snyk · Prowler

PROTECT

Stop triaging vulnerability notifications one finding at a time

Sysdig's preview automations collapse Vulnerability Findings notifications into one per resource or image rather than one per finding, which is the difference between a ticket queue and a channel nobody reads. Checkmarx's Fusion SAST ships with Triage, Remediation and Developer

> Assist to push the same compression into the fix step, generating patches and surfacing them in the IDE.
>
> Sysdig Secure

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

nerva

Feedback ⓘ v1.84.0 NOTES ↗

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

└──▷ GET THIS VERSION ▾

Go

```
$ go install github.com/praetorian-inc/nerva/cmd/nerva@v1.84.0
```

git

```
$ git clone --branch v1.84.0 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.84.0
```

Go module path matches the repo · checked 30 Oct Wrong command?

NEW DETECTION VMware management interface exposure detection

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Semgrep

Feedback ⓘ changelog · Oct 30, 2026 CHANGELOG ↗ RANK

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

> Semgrep's latest release adds a new related-issues API endpoint and direct Anthropic AI provider connectivity, expands Supply Chain dependency coverage and filtering, adds PowerShell support to Guardian, and ships several breaking changes including removal of QL language support and a renamed xAI provider.

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

01 **Related issues API endpoint for findings** NEW 80

A new experimental endpoint, `GET /api/deployments/{deployment_id}/issues/{issue_id}/related-issues`, lets you query a finding's related issues directly. The existing `relatedIssues` field on finding responses is now deprecated in favor of this endpoint.

Retrieve all related issues for a specific finding to triage clusters of related vulnerabilities at once.

```
$ curl -H 'Authorization: Bearer <token>'
'https://semgrep.dev/api/deployments/<deployment_id>/issues/<issue_id>/related-issues'
```

— Named endpoint with method, path and runnable curl example changelog-20261030-70018976

02 **QL language support removed** BREAKING 75

Support for QL (the CodeQL query language) is removed: `--lang ql` and the `ql` language key in rules are now rejected, and `.ql` and `.qll` files are no longer parsed.

— Names flag, rule key and file extensions affected changelog-20261030-70018976

03 **Full requirements.txt coverage in Supply Chain** IMPROVED 65

Semgrep Supply Chain now includes every package declared in a Python `requirements.txt` file — including packages with version ranges (e.g. `requests>=2.0`) and packages without a version (e.g. `flask`) — in the dependency inventory and SBOM, not only packages pinned with `==`.

— Clear before/after with concrete package examples changelog-20261030-70018976

04 **AI provider settings: Azure OpenAI Base model and default provider change**

65

BREAKING

AI provider settings now include a required `Base model` field for Azure OpenAI deployments, needed for new Azure OpenAI setups. Organizations that had not chosen an AI provider mode now default to the `Cloudflare Workers AI` mode, while organizations that chose their own providers or are on the minimal data retention policy remain unchanged.

— Names exact field, default mode and who is affected `changelog-20261030-70018976`

05 **xAI provider renamed to SpaceXAI** BREAKING

60

The `xAI` AI provider is renamed to `SpaceXAI` and its endpoint field is removed; the provider always uses `https://api.x.ai/v1`.

— Names old/new provider and fixed endpoint URL `changelog-20261030-70018976`

06 **Rate limit change for public deployment endpoints** BREAKING

60

Rate limits for `/api/v2/deployments/public/{id}` and `/api/v2/deployments/public/search` now count every request including repeated and concurrent requests.

— Names exact endpoints affected by the stricter counting `changelog-20261030-70018976`

THINNER COVERAGE BELOW

07 **Advisories page filters and Incident column** IMPROVED

55

The Advisories page gains Language, Severity, and Advisory type filters, matching the filters already available on the Findings page, plus a new Incident column in the advisories table.

— Names filters and column; clear UI starting point `changelog-20261030-70018976`

08 **Quick dependency version filtering** IMPROVED

55

Dependencies can now be filtered by version directly on the Dependencies page without opening advanced options; the dependency filter is now an `All dependencies` menu with `Exact` and `Contains matching` modes and `All versions` or `Select versions` options.

— Names exact menu and modes; UI path given `changelog-20261030-70018976`

09 **CSV export for Agentic Workflows issues** NEW

55

Selected Agentic Workflows issues can now be exported to CSV; Semgrep exports only the issues or groups selected, or uses active filters when nothing is selected, with a 10,000-issue limit applied to the selection.

— Names export format, limit and selection logic changelog-20261030-70018976

10 **False positive rate column in Policies rules table** NEW 55

A sortable False positive rate column is added to the rules table in Policies, covering Semgrep Code and Semgrep Secrets rules, with the ability to select a rate to open associated false-positive findings.

— Names exact column and linked rule types with interaction changelog-20261030-70018976

11 **Direct Anthropic AI provider connection** NEW 50

Users can connect directly to Anthropic by supplying their own Anthropic API key under the custom AI provider mode in Semgrep AppSec Platform settings, for use with in-app features. Available to Semgrep Enterprise and existing bring-your-own-key users.

— Names setting location and eligibility but no exact config key changelog-20261030-70018976

12 **Finding type filter for dashboards** NEW 45

A Finding type filter is added to dashboards to view AI-powered and Rule-based findings separately in charts, drilldowns, and PDF exports.

— Names the filter and its views but little mechanism changelog-20261030-70018976

13 **PowerShell support in Semgrep Guardian** NEW 35

Semgrep Guardian now supports PowerShell, running its hooks on both PowerShell commands and Bash commands.

— Brief statement with no command detail or scope changelog-20261030-70018976

└──▷ **BREAKING ON UPGRADE**

! Support for QL (the CodeQL query language) is removed: `--lang ql` and the `ql` language key in rules are now rejected, and `.ql` and `.qll` files are no longer parsed.

! The `xAI` AI provider is renamed to `SpaceXAI` and its endpoint field is removed; the provider always uses `https://api.x.ai/v1`.

! Azure OpenAI settings now require a Base model field; new Azure OpenAI setups must supply this value.

! Rate limits for `/api/v2/deployments/public/{id}` and `/api/v2/deployments/public/search` now count every request including repeated and concurrent requests.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Integration updates (2026-10-30) BLOG ↗

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

Veracode integrations add IaC policy and predependency command parameters to veracode.yml across GitHub, GitLab, and Azure DevOps workflows, environment setup checks in VS Code, and platform-managed GitHub install and scan scheduling.

Add a `predependency_command` parameter to `veracode.yml` in the GitHub, GitLab, and Azure DevOps Workflow Integrations to run commands before Veracode packages the project.

NEW INTEGRATION Add a `build_runs_on` parameter to `veracode.yml` in the GitHub Workflow Integration for SCA scans.

Add a `policy` parameter to `veracode.yml` in the GitHub, GitLab, and Azure DevOps Workflow Integrations for IaC scans.

Add a custom agent pool name parameter to `veracode.yml` in the Azure DevOps Workflow Integration.

Veracode Scan for VS Code now adds environment setup checks to identify connectivity, proxy, DNS, firewall, and TLS issues.

+6 more

NEW SUPPORTED PLATFORM JetBrains IDE versions 2024.2 through 2026.3

NEW SUPPORTED PLATFORM Jira 11 support in Veracode Integration for Jira Server

UPDATED INTEGRATION Autopackaging support for Apex, VB6, and Perl in GitHub Workflow Integration

└──▷ **BREAKING ON UPGRADE**

! Veracode Azure DevOps Extension v4.0.0 deprecates Veracode Flaw Importer v3 and replaces it with VeracodeFlawImporter v4; existing pipelines must be manually migrated from v3 to v4.

! Java API wrapper v26.10.20.1 now updates the repository URL in the Veracode application profile only when `-repourl` is explicitly specified.

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ Webhook IP ranges added to the Meta API (2026-10-29)

BLOG ↗ RANK

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

> GitHub expanded its Meta API with new webhook IP ranges that will break static allowlists if not updated, and brought Copilot plugin management to general availability with repository-level defaults and enterprise-wide controls.

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

01 New webhook IP ranges in Meta API

BREAKING

85

The Meta API's `hooks` field now includes three new CIDR prefixes — `9.234.98.160/28`, `9.234.98.176/28`, and `9.234.106.48/28` — covering all github.com plans including GitHub Enterprise Cloud. Static network allowlists for incoming webhook deliveries must be updated to include these ranges or they will reject the new traffic.

Pull the current webhook IP ranges from the Meta API to update a static network allowlist before the new CIDRs go live.

```
$ curl -s https://api.github.com/meta | jq '.hooks'
```

— Exact CIDRs, affected field, and a runnable query to fetch them.

Webhook IP ranges added to the Meta API…

02 Copilot plugin defaults and enterprise plugin controls

NEW

60

Repository maintainers can now define default plugins and a plugin marketplace so contributors get consistent Copilot plugin setup without repeating configuration per project, and contributors can turn off those repository-default plugins where enterprise policy permits. Enterprise administrators can require plugins (auto-installed and user-locked), block plugins, and restrict installation to approved marketplaces through managed settings that override repository and personal choices. The controls apply across VS Code, Copilot CLI, and the GitHub Copilot app, with managed plugin controls available on GitHub Enterprise Cloud for Copilot Business and Copilot Enterprise licenses, and repository defaults also available to individual Copilot users.

— Describes override hierarchy and license scope but no concrete config keys or commands.

Repository plugin defaults and enterpri…

└──▷ **BREAKING ON UPGRADE**

! Static network allowlists for incoming webhook deliveries will reject traffic from `9.234.98.160/28`, `9.234.98.176/28`, and `9.234.106.48/28` unless those CIDR prefixes are added manually.

WAS THIS USEFUL? 👍 👎

qodana-cli

Feedback ⓘ v2026.2.3 NOTES ↗ RANK

JetBrains Qodana's official command line tool

> qodana-cli adds automatic exchange of an organization API token for a project-scoped token, narrowing the credential scope needed for CI scans.

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

01 **Automatic org-to-project API token exchange** IMPROVED

33

qodana-cli now automatically exchanges an organization API token for a project-scoped token, reducing the credential scope needed for CI scans.

— No mechanism details, flag, or config key named. v2026.2.3

WAS THIS USEFUL? 👍 👎

Checkmarx

Feedback ⓘ

What the PHP Patch Scramble Reveals About Remediation (2026-10-29) BLOG ↗

RANK

Checkmarx agentic application security software platform Combines Hybrid scanning, AI-powered agents, and unified risk intelligence across every attack surface

> Checkmarx introduces Fusion SAST, an AI-augmented scanning engine, alongside a trio of AI Assist capabilities — Triage Assist, Remediation Assist, and Developer Assist — that automate finding triage, fix generation, and in-IDE remediation.

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

01 **Remediation Assist with scoped automation and project overrides**

NEW 85

Remediation Assist automatically generates AI remediation suggestions after a scan or automated triage completes, with admins defining the projects, branches, scanners, and severity thresholds in scope. Project overrides let individual teams adapt these automated remediation settings where permitted.

After enabling automated remediation, configure scope so fixes are queued the moment a scan finishes — no manual request per finding.

> 📍 In Checkmarx One, go to Settings › Automated AI Remediation, set the target projects, branches, scanners, and severity thresholds, then enable Generate fix after scan completes.

— Includes exact navigation steps and config scope fields

What the PHP Patch Scramble Reveals Abo…

02 **Fusion SAST with Findings Analysis Engine** NEW 65

Fusion SAST combines established rule-based security analysis with frontier-model reasoning, and a Findings Analysis Engine consolidates results to reduce false positives and duplicate findings. Checkmarx reports an F1 score of 0.74 versus 0.64 for the Next-Gen SAST foundation.

— Names mechanism and benchmark numbers but no config/usage path

What the PHP Patch Scramble Reveals Abo…

THINNER COVERAGE BELOW

03 **Triage Assist for automated finding exploitability review**

NEW

55

Triage Assist (AI Triage) evaluates whether a finding is reachable and exploitable, automatically setting a finding's state to Confirmed or Proposed Not Exploitable, with its reasoning visible in Risk Orchestration.

– Names exact state values and visibility location but no steps to enable

What the PHP Patch Scramble Reveals Abo…

04 **Developer Assist in IDE and AI coding workflows** NEW

40

Developer Assist brings security analysis and remediation guidance into IDE and AI coding workflows, including supported autonomous workflows where it works with the coding agent to repeat remediation and rescanning.

– Describes capability in prose without named flags or exact setup

What the PHP Patch Scramble Reveals Abo…

WAS THIS USEFUL? 👍 👎

Snyk Open Source

Feedback ⓘ changelog · Oct 29, 2026 BLOG ↗

Snyk Open Source is a developer security tool that scans project dependencies for known vulnerabilities and license issues, suggesting fixes or upgrades to remediate them.

NEW SUPPORTED PLATFORM Python 3.15 support

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Oct 29, 2026 BLOG ↗ RANK

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

> Snyk API & Web now extends its security header, cookie attribute, and SSL/TLS checks to every Extra Host configured on a target, not just the primary target URL.

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

01 **Security checks extended to Extra Hosts** BREAKING 57

Snyk API & Web now applies security header, cookie attribute, and SSL/TLS checks — including HSTS header not enforced and Missing Content Security Policy header — to every Extra Host configured on a target, not just the primary target URL endpoint. As a result, scans of targets with Extra Hosts configured may report new findings for those hosts on their next scan even if nothing on the target itself has changed.

— Names exact checks and scope change, but no config surface to act on.

2026-10-29-0107292c

└──▷ **BREAKING ON UPGRADE**

! Scans of targets with Extra Hosts configured will now report new security header, cookie attribute, and SSL/TLS findings for those hosts — existing targets may show new findings on their next scan even if nothing on the target has changed.

WAS THIS USEFUL? 👍 👎

GitLab

Feedback ⓘ

GitLab Orbit is GA: Higher agent accuracy, lower token cost (2026-10-29) BLOG ↗

RANK

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

> GitLab's Orbit context graph reaches general availability for Premium and Ultimate customers, giving agents and engineers a live, queryable graph of code, pipelines, vulnerabilities and ownership — paired with new ways to query it, from CLI-wired MCP agents to plain-language vulnerability blast-radius hunts and exact regex code search.

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

01 Agent and API access to the Orbit graph NEW 85

MCP-compatible agents (Claude Code, Codex, Cursor, OpenCode, Gemini CLI, Antigravity) connect to Orbit by running `glab auth login` then `glab orbit setup`; custom agents and internal tooling connect through REST; agents on GitLab's Duo Agent Platform query Orbit natively with nothing to configure; and engineers can query the graph directly through the Data Explorer.

Wire up an MCP-compatible agent such as Claude Code or Cursor to Orbit so it can query live codebase context instead of crawling files.

```
$ glab auth login
glab orbit setup
```

— Exact CLI commands given; REST and Data Explorer access described only in prose

GitLab Orbit is GA: Higher agent accura… GitLab Orbit: Find vulnerability blast …

02 Orbit live context graph reaches GA NEW 80

GitLab Orbit, a live queryable graph connecting code, merge requests, pipelines, vulnerabilities, work items, and ownership across all projects in a group, is now generally available for Premium and Ultimate customers on `GitLab.com`, with Self-Managed and GitLab Dedicated support coming in the following weeks and a free 30-day trial available through account teams. Orbit indexes every push, merge, and deploy in real time — at GitLab scale, more than 1.5 million projects and 8 TB of graph data in under 2 hours — so agents reason from current lifecycle state. Orbit access follows existing GitLab roles so people and agents see only what they are permitted to see, and the graph queries behind each answer can be audited. Orbit is sold as query packs (a

shared pool per org) priced at $30 in GitLab Credits per user per month for 1,000 queries, with overages billed at 33 queries per GitLab Credit (~$0.03 per query).

— Scale, pricing and GA scope named; no runnable command for the core graph itself

GitLab Orbit is GA: Higher agent accura… | GitLab Orbit: Find vulnerability blast …

---

03 **Vulnerability blast radius queries via Duo Agentic Chat** NEW

80

Security teams can ask GitLab Duo Agentic Chat plain-language questions against the Orbit graph to find blast radius without manual correlation, for example: 'Use Orbit to map the blast radius of PHP in this group. List every project that contains PHP code, grouped by subgroup, with owners.' GitLab's built-in Security Analyst Agent now uses Orbit directly to traverse the graph and answer these questions, and teams can also collect every container scanning finding for specific CVEs across a group via Agentic Chat (requires Ultimate and Security Manager role).

Map every project in a group that uses PHP, grouped by subgroup with owners named, during a vulnerability triage.

> 📍 In GitLab Duo Agentic Chat, type: Use Orbit to map the blast radius of PHP in this group. List every project that contains PHP code, grouped by subgroup, with owners.

Collect every container scanning finding for a specific CVE across the group via Agentic Chat (requires Ultimate and Security Manager role).

> 📍 In GitLab Duo Agentic Chat, type: Collect every container scanning finding for CVE-2026-109618 and CVE-2026-109619 across the group from the vulnerability report.

— Concrete chat queries and role requirement given; no API form shown

GitLab Orbit: Find vulnerability blast …

---

04 **Exact regex code search across a group** NEW

75

At the top-level group, exact code search with regular expression mode lets engineers find things like container base images alongside Orbit graph queries, e.g. `lang:Dockerfile ^FROM.*php` or `file:composer\.json$`.

Find all Dockerfiles that pull a PHP base image across the entire top-level group using exact regex search alongside Orbit.

```
$ lang:Dockerfile ^FROM.*php
```

— Exact search syntax given; scope and limits of the regex engine not detailed

GitLab Orbit: Find vulnerability blast …

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

kingfisher

Feedback ⓘ v2.16.0 NOTES ↗ RANK

Detect secrets + live validation. Map blast radius.

> Kingfisher's v2.16.0 release expands its built-in rule catalog to 501 rules with new connection-string password detectors, adds detection of bare AWS access IDs, and updates its Betterleaks integration with new content detectors.

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

01 Connection-string password rules added to catalog NEW

68

Five new connection-string password rules now ship by default, covering PostgreSQL/MySQL/MariaDB JDBC query passwords and SQL Server JDBC properties. They are embedded by default in the CLI, `kingfisher-rules`, and the Python SDK, bringing the total rule catalog to 501 rules.

— Names rule types, surfaces, and exact catalog count. v2.16.0

THINNER COVERAGE BELOW

02 AWS access ID detection without nearby secret key NEW

45

Kingfisher now detects AWS access IDs even when no secret key is found nearby, though validation and revocation still require a matching secret key.

— Explains behavior change but no config or flag given. v2.16.0

03 Betterleaks integration updated with new detectors IMPROVED

40

The Betterleaks integration now reflects v2.0.0-rc.5, adding eight new content detectors for Labelbox, NuGet, and Tailscale.

— Names version and providers but no usage detail. v2.16.0

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS

Feedback ⓘ

Sync your GitHub teams, members, and perimeters automatically (2026-10-29)

BLOG ↗ 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 adds automatic GitHub team sync into workspaces, extends custom webhooks to support multiple headers, ships new Public API endpoints for incident/ticket management, tightens Endpoint Protection scope grants in ggshield, and broadens Bridge routing for host-derived secrets.

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

01 **Automatic GitHub team sync** NEW 80

Enabling team synchronization on a GitHub integration makes GitGuardian mirror the GitHub organization's teams, members, and repository ownership into the workspace. Sync runs automatically every 24 hours or on demand via Sync now, with results shown in a synchronization history view.

Trigger or review GitHub team sync after enabling it, to immediately mirror your GitHub org's teams and members without waiting for the 24-hour cycle.

📍 Open your GitHub integration settings, scroll to Team synchronization, click Add organization, select the organization to mirror, then click Sync now to run an immediate sync and view results in the synchronization history.

— Exact UI steps given; no API endpoint named Sync your GitHub teams, members, and pe…

02 **Opt-in Endpoint Protection scopes in ggshield auth login** 80

BREAKING

The `ggshield auth login` authorization page now groups Endpoint Protection scopes into a single opt-in card and no longer grants them by default; existing automation that relied on those scopes being pre-selected will need to re-authenticate.

— Exact command named with clear migration note Sync your GitHub teams, members, and pe…

03 **Bridge routing for host-derived secret detectors** IMPROVED 60

GitGuardian Bridge now routes validity checks and secret analysis for detectors where the host is read from the secret itself (such as database connection strings) whenever

that host is mapped to a bridge — previously this routing was limited to detectors with a custom host configured.

— Clear before/after but no user-facing config action

Sync your GitHub teams, members, and pe…

THINNER COVERAGE BELOW

04 **Multi-header custom webhooks** IMPROVED 50

Custom webhook integrations now accept multiple custom headers per webhook — such as bearer tokens or routing headers — configured from the webhook's settings page.

— Names UI location but no exact config key Sync your GitHub teams, members, and pe…

05 **Public API endpoints for incident tickets and occurrence content**

NEW 45

New Public API endpoints let you unlink a ticket from an incident, and retrieve the stored content of an occurrence of an internal or public incident in clear text — the latter subject to the same access requirements as secret access.

— Endpoints described without exact path or method

Sync your GitHub teams, members, and pe…

└──▷ **BREAKING ON UPGRADE**

! Endpoint Protection scopes are no longer granted by default during `ggshield auth login` — users must explicitly opt in via the new single card, so existing automation that relied on those scopes being pre-selected will need to re-authenticate.

WAS THIS USEFUL? 👍 👎

GitGuardian Detection Engine

Feedback ⓘ

Detection Engine Updates Version 2.175 (2026-10-27) BLOG ↗

GitGuardian Detection Engine scans source code repositories, commits, and pipelines to detect leaked secrets such as API keys, passwords, and credentials.

NEW DETECTION Portainer API Key detector

NEW DETECTION Weights & Biases API Key checker and analyzer

NEW DETECTION Custom-host checkers activated for eight services

UPDATED DETECTION json_web_token recall improvement

UPDATED DETECTION sentry_org_auth_token recall improvement

UPDATED DETECTION jenkins_api_token host URL matching

UPDATED DETECTION keycloak_access_token precision improvement

UPDATED DETECTION generic_password false positive reduction

UPDATED DETECTION generic_high_entropy_secret false positive reduction

UPDATED DETECTION microsoft_azure_storage_account_key_with_account_name moved to own detector group

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

zarf

Feedback ⓘ v0.89.0 NOTES ↗ RANK

The Airgap Native Package Manager for Kubernetes

> Zarf v0.89.0 introduces the v1beta1 package format with major schema changes — granular SBOM/attachment layers, an object-based `differential` field, and removed `###ZARF` builtins — alongside custom OCI credential providers, granular layer selection in the `zoci` Load API, and several breaking changes including opt-in kustomize plugins, a removed `cachepath` SDK option, and removed OCI reference shortcut aliases.

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

01 v1beta1 package schema and init format

BREAKING

85

`v1beta1` introduces a new init package format for cluster initialization, splits SBOM and documentation into separate granular files and layers (the `documentation` field is renamed to `attachments`), changes `differential` from a scalar to an object, improves component config override import logic for more predictable composition, removes `###ZARF` builtin substitutions in favor of explicit values, adds an `id` field to the image object in the package schema, and validates annotation size and package secret size at `zarf package create` time to reject oversized values early; package data is no longer duplicated in the package secret.

— Names every changed field and config surface; no exact CLI example given v0.89.0

02 Granular layer selection in zoci Load API

BREAKING

65

zarf now supports granular layer selection from the `zoci` Load API, letting callers fetch only the layers they need instead of requiring all layers; existing code relying on the previous all-layers default will break.

— Names the API and behaviour change clearly, no code sample v0.89.0

03 Per-operation cache configuration replaces cachepath

BREAKING

60

The global `cachepath` SDK option is removed; cache paths must now be set per-operation, and image layer caching is simpler and no longer shared across concurrent operations.

— Names the removed SDK option and describes new caching behaviour v0.89.0

THINNER COVERAGE BELOW

04 **Kustomize plugin opt-in for dev commands** BREAKING 50

Kustomize plugins are no longer active by default on read-only dev commands; users must now explicitly opt in via a flag rather than having them active by default.

— Describes behaviour change but flag name not given v0.89.0

05 **Custom credential providers for OCI registries** NEW 45

zarf now supports custom credential providers for OCI registries, enabling non-standard auth flows when pushing or pulling packages.

— Explains purpose but no config key or flag given v0.89.0

06 **Removal of OCI reference shortcut aliases** BREAKING 40

Defense-unicorn shortcut aliases (short-form OCI refs) are removed from OCI package references; full references are now required.

— States the removal and requirement, limited further detail v0.89.0

07 **Explicit injector image required in proxy mode** BREAKING 40

Proxy mode now requires an injector image to be specified explicitly; deployments that omitted it will fail.

— States requirement and failure mode but no config key shown v0.89.0

08 **Release candidate versions pass version checks** IMPROVED 20

Release candidate versions now pass `zarf` version requirement checks.

— Minimal description, no mechanism or example v0.89.0

└──▷ BREAKING ON UPGRADE

! Kustomize plugins are no longer active by default on read-only dev commands; you must now explicitly opt in.

! The `zoci` Load API now requires callers to specify granular layers; existing code that relied on the previous all-layers default will break.

! Package creation now fails if the package secret exceeds the size limit (`zarf package create`).

! In `v1beta1` , `differential` is now an object; configs using it as a scalar will be rejected.

- ! In `v1beta1`, the `documentation` field is renamed to `attachments`; any package definition using `documentation` must be updated.
- ! In `v1beta1`, annotations exceeding the new size limit will be rejected at create time.
- ! The `id` field is added to the image object; tools parsing the image schema must handle the new field.
- ! `###ZARF` builtin substitutions are removed from `v1beta1`; packages relying on them must be rewritten to use explicit values.
- ! Defense-unicorn shortcut aliases (e.g. short-form OCI refs) are removed from OCI package references; full references are now required.
- ! In `v1beta1`, SBOM and documentation are now stored as separate granular files and layers; tooling that reads the old combined format will break.
- ! Package data is no longer duplicated in the package secret; code that reads package data from the secret will stop working.
- ! Image layer caching is now simpler and no longer shared across concurrent operations; workflows relying on shared cache behaviour will need review.
- ! Proxy mode now requires an injector image to be specified explicitly; deployments that omitted it will fail.
- ! The global `cachepath` SDK option is removed; any SDK consumer setting it must switch to per-operation cache configuration.

WAS THIS USEFUL? 👍 👎

Endor Labs

Feedback ⓘ

Endor Labs + Google Antigravity: Secure code by default in agentic workflows | Endor Labs

BLOG ↗ RANK

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

> Endor Labs shipped an Antigravity Plugin Marketplace integration that brings its MCP server and AURI Agents into Google Antigravity's agentic development workflows, adding real-time dependency and secret scanning plus automated remediation, decision-support, and vulnerability-explanation agents.

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

01 SCA Remediation agent for automated dependency fixes NEW

90

The SCA Remediation AURI Agent runs inside Antigravity on the developer's own infrastructure using their own LLM keys and source-control credentials. Triggered with a prompt such as 'Use the SCA Remediation agent to fix the vulnerable dependency in this project,' it applies reachability and upgrade impact analysis, selects the safest upgrade, opens a pull request, and attaches evidence to the PR body.

Trigger an automated dependency fix inside an Antigravity session — the SCA Remediation agent selects the safest upgrade and opens a PR with evidence.

> 📍 Steps: In Google Antigravity, open the Plugin marketplace and install the Endor Labs plugin. Then in your Antigravity session, ask: 'Use the SCA Remediation agent to fix the vulnerable dependency in this project.' The agent applies reachability and upgrade impact analysis, opens a pull request, and attaches evidence to the PR body.

— Exact trigger prompt and end-to-end mechanism given

Endor Labs + Google Antigravity: Secure…

02 Vulnerability Explainer agent NEW

75

The Vulnerability Explainer AURI Agent turns a CVE ID into a plain-language explanation of impact and remediation steps specific to the project. Asking 'Use the Vulnerability Explainer agent to explain CVE-2024-12345 in the context of this project' returns a project-specific impact summary and recommended remediation steps.

Get a plain-language explanation of what a CVE means for your specific project before deciding whether to act on it.

📍 Steps: With the Endor Labs plugin installed in Antigravity, ask: 'Use the Vulnerability Explainer agent to explain CVE-2024-12345 in the context of this project.' The agent returns a project-specific impact summary and recommended remediation steps.

— Concrete prompt example shows exact usage Endor Labs + Google Antigravity: Secure…

03 **Antigravity Plugin Marketplace integration with MCP server**

NEW 67

Endor Labs added a one-click install from Google Antigravity's Plugin marketplace that brings the Endor Labs MCP server and AURI Agents (including AI SAST Triage) into Antigravity-based agentic workflows. The MCP server gives Antigravity agents real-time security context — checking packages for known vulnerabilities and malicious code before a dependency is added, scanning for secrets, and surfacing vulnerability details inline while code is written.

— Names MCP server and agents but install flow only described, not a command

Endor Labs + Google Antigravity: Secure…

THINNER COVERAGE BELOW

04 **Dependency Decision Helper agent** NEW 40

The Dependency Decision Helper AURI Agent answers whether a package should be added before an agent installs it, running inside Antigravity on the developer's own infrastructure.

— Purpose stated but no example or trigger shown Endor Labs + Google Antigravity: Secure…

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ

Harden Runner Releases: October 29, 2026 (2026-10-29) BLOG ↗ RANK

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's Harden Runner adds a flag to block implicit AWS endpoint allowances on self-hosted Linux runners, exposes new eBPF monitoring health metrics for ARC-hosted runners, and extends support to Avrea-hosted enterprise runners.

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

01 eBPF monitoring health metrics and configurable buffers in ARC Harden Runner

NEW 90

ARC Harden Runner v2.24.0 reports eBPF monitoring health metrics — ring buffer usage, dropped network/DNS/process events, and DNS tracking map inserts, removes, and errors — via a Prometheus endpoint when `env.enableMetrics` is enabled. Buffer sizing is configurable via keys such as `dnsTrackingMapSize` and `processRingBufferSize`.

Tune buffer sizes on a high-throughput ARC node where DNS lookups or process events are being dropped, and confirm the sizing is working via Prometheus metrics.

yaml

```yaml
env:
  dnsTrackingMapSize: 65536
  processRingBufferSize: 32768
  enableMetrics: true
```

— Names config keys, metrics tracked, and the Prometheus export mechanism.

Harden Runner Releases: October 29, 202…

02 Block implicit AWS endpoint allowances on self-hosted runners

NEW 80

A new `disable_implicit_endpoints` feature flag stops self-hosted Linux runners from implicitly allowing AWS endpoints, while GitHub and StepSecurity endpoints remain allowed. The flag shipped in both Self-Hosted VM Harden Runner v1.9.5 and GitHub Action v2.22.3 (enterprise tier).

Prevent the agent from silently opening AWS endpoints on a self-hosted Linux runner — useful when your policy requires all allowed endpoints to be explicitly declared.

```yaml
env:
  disable_implicit_endpoints: true
```

— Names exact flag and two shipping products with a runnable config example.

Harden Runner Releases: October 29, 202…

THINNER COVERAGE BELOW

03 **Avrea-hosted runner support** NEW 30

Harden Runner now supports runners hosted by Avrea (enterprise tier), added in GitHub Action v2.22.3.

— Names the provider and tier but no setup detail given.

Harden Runner Releases: October 29, 202…

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Invicti

Feedback ⓘ Release 20261028 (2026-10-28) BLOG ↗ RANK

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

> Invicti's latest release adds client certificate secrets for target authentication, enriches the API catalog with sensitive-data labels and scanning filters plus discovery notifications, and introduces a Scan CLI environment variable to halt DAST scans on policy failure.

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

01 **API catalog sensitive-data labels, scanning filter, and discovery notifications**

NEW 65

The API catalog and API discovery now label each API and endpoint with the category of sensitive data it exposes — Secrets, PII, or Financials — so risk is visible without opening each finding. A new Requires scanning filter surfaces newly discovered or recently updated APIs that still need a DAST scan. Users also receive bell notifications and emails when APIs are discovered or updated, with a direct link to the changed APIs in the catalog; email alerts can be turned off in profile notification settings.

– Names labels, filter, and notification controls with opt-out location

Release 20261028 (2026-10-28)

02 **INVICTI_FAIL_FAST environment variable for Scan CLI** NEW

60

A new `INVICTI_FAIL_FAST` environment variable stops DAST scans on first policy failure when run via the Scan CLI.

– Names exact env var but lacks usage example or further mechanism

Release 20261028 (2026-10-28)

THINNER COVERAGE BELOW

03 **Client certificate secrets for target authentication** NEW

50

Client certificate files can now be stored as secrets and selected in a target's client certificate authentication settings, rather than being configured per-target directly.

– Describes the workflow but no config key or API detail Release 20261028 (2026-10-28)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL?

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

cortex-analyzers

Feedback ⓘ 3.10.2 NOTES ↗ RANK

Cortex Analyzers Repository

> Cortex-Analyzers 3.10.2 enhances the ANY.RUN sandbox analyzer with proxy and domain flexibility, and adds artifact support to the analyzer testing framework.

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

01 **Proxy and automated interactivity in ANY.RUN analyzer**

IMPROVED 60

The ANY.RUN analyzer now supports proxy configuration, alternative domains, and `opt_automated_interactivity` for sandbox submissions.

– Names config options but no usage example or limits given 3.10.2

THINNER COVERAGE BELOW

02 **Artifact support in TestAnalyzer** NEW 35

TestAnalyzer now supports artifacts, allowing artifact-based observable testing during analyzer development.

– States capability but no concrete usage detail 3.10.2

WAS THIS USEFUL?

NETWORK SECURITY

◆ OT / ICS Security

Xage Security

Feedback ⓘ announced · Oct 30, 2026 BLOG ↗

Xage Security is a company providing zero-trust access control and data protection for industrial operational technology and critical infrastructure, enabling secure remote access, identity-based policy enforcement, and encrypted data sharing.

Xage adds a SailPoint Identity Security integration for bidirectional sync of User and User Group entitlements, automated provisioning, and revocation.

NEW INTEGRATION Xage now adds a SailPoint Identity Security integration that lets SailPoint provision, update, disable, or remove Xage Users and User Group memberships through governed lifecycle events, requests, and approvals.

Xage returns current users and User Group memberships to SailPoint for access discovery, certification, reconciliation, and remediation, closing the loop between governance decisions and enforcement state.

Automated revocation now removes the corresponding Xage User Group membership when SailPoint revokes an entitlement, without requiring a separate manual cleanup step.

WAS THIS USEFUL? 👍 👎

cisagov Malcolm

Feedback ⓘ v26.10.0 NOTES ↗ RANK

Malcolm is a powerful, easily deployable network traffic analysis tool suite for full packet capture artifacts (PCAP files), Zeek logs and Suricata alerts.

> Malcolm v26.10.0 adds normalized ECS threat-intelligence fields across every intel source, Community ID computation in Logstash, expanded OT protocol port and field mapping, and a richer document-lookup API.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 ZEEK_PORTS_ environment variables for OT analyzers

BREAKING

90

New `ZEEK_PORTS_…` environment variables register non-default ports for thirteen ICSNPP OT protocol analyzers (BACnet, C12.22, EtherNet/IP, GE-SRTP, Genisys, HART-IP, Modbus, Omron FINS, OPC UA Binary, Profinet IO-CM, ROC+, S7comm, and Synchrophasor), so a site running these protocols on nonstandard ports no longer needs a custom Zeek script. As part of this, the older `ZEEK_SYNCHROPHASOR_PORTS`, `ZEEK_GENISYS_PORTS`, and `ZEEK_ENIP_PORTS` variables are renamed to `ZEEK_PORTS_SYNCHROPHASOR`, `ZEEK_PORTS_GENISYS`, and `ZEEK_PORTS_ENIP`; existing values are migrated on upgrade, but external automation referencing the old names must be updated.

Run Modbus on a nonstandard port (502 is occupied; site uses 5020) so Malcolm parses it without a custom Zeek script.

```env
ZEEK_PORTS_MODBUS=5020/tcp
```

— Names all 13 protocols, rename mapping, and a runnable config example v26.10.0

02 Normalized ECS threat.* enrichment from all intel sources

NEW

85

Threat intelligence enrichment now maps into normalized ECS `threat.*` fields across every source: Zeek intel kill chain phases map to `threat.tactic.*`; MISP MITRE galaxy tags in `cif_tags` map to `threat.technique.*` and `threat.technique.subtechnique.*`; Mandiant and Google Threat Intelligence actors map to `threat.group.name`; and Suricata Emerging Threats `mitre_*` metadata and capa subtechnique results are mapped the same way, with ATT&CK for ICS distinguished from Enterprise. Per-indicator metadata is also carried into Zeek intel, including Mandiant categories, campaigns, reports, opt-in threat rating

(`meta.threat_score`), and STIX 2.1 `confidence` values, with Mandiant's M-Score treated as confidence and absent values left unset.

— Deep mechanism and many named fields but no runnable command v26.10.0

---

03 **Community ID computation in Logstash for Zeek logs** NEW

85

Logstash can now optionally compute Community ID for Zeek logs that arrive without one, controlled by the `LOGSTASH_ZEEK_COMPUTE_COMMUNITY_ID` environment variable, enabled by default. An ID Zeek already supplied is never overwritten, and transport is inferred by log type for single-transport analyzers.

Enable Community ID computation for Zeek logs that arrive in Malcolm without one, so correlation across tools works even when Zeek did not tag the flow.

```env
LOGSTASH_ZEEK_COMPUTE_COMMUNITY_ID=true
```

— Exact env var plus runnable example and overwrite behavior v26.10.0

---

04 **Malcolm API document-lookup arguments and status fields**

IMPROVED

85

The Malcolm API's document lookup gains `sort` , `offset` , and `track_total_hits` arguments and a `total` hit count with its exact-or-lower-bound relation; both the document lookup and field aggregations endpoints now return `shards` and `timed_out` status so a client can distinguish a partial or timed-out search from a genuinely empty result. Argument validation errors return HTTP 400 with an `error` string, and setting `MALCOLM_API_DEBUG=true` logs a compact per-request search summary.

— Named params and env var but no concrete call example v26.10.0

---

05 **OT protocol field normalization to Arkime names** IMPROVED

70

Zeek OT protocol fields are normalized to Arkime field names for BACnet, C12.22, DNP3, EtherNet/IP (ENIP/CIP), Modbus, OPC UA, S7comm, and Synchrophasor so filters and drill-downs carry across Dashboards and Arkime; original `zeek.*` fields are preserved alongside the new top-level fields.

— Names protocols and preservation behavior, no direct command v26.10.0

---

06 **Zeek 9.0.1 LTS with multicast and NTP log support** NEW

65

Zeek is updated to 9.0.1 LTS, parsing and indexing two new log types: `multicast_participants.log` , surfaced in a new Multicast dashboard, and

`ntp_control.log` and `ntp_private.log`, added to the index template and Arkime field definitions.

— Names log files and dashboard but only a UI starting point v26.10.0

THINNER COVERAGE BELOW

07 **Hedgehog Linux sensor kiosk startup prompt** NEW 55

Hedgehog Linux ISO installation now prompts whether to run the sensor status kiosk at startup, defaulting to yes; choosing `no` disables `kiosk.service` and removes the browser autostart entry.

— Names the service and default, but no command example v26.10.0

08 **Google Threat Intelligence collection filtering parameters** NEW 50

New `ioc_types` and `include_autogenerated_tags` parameters for Google Threat Intelligence collections let a deployment load only the indicator types it needs and omit machine-generated tags, reducing memory cost in Zeek.

— Names two parameters and their purpose, no usage example v26.10.0

09 **NetBox device-type import vendor filtering** NEW 50

`NETBOX_DEVICETYPE_IMPORT_VENDORS` limits the NetBox device-type library import to selected manufacturers on first import, also constraining which manufacturers are matched during MAC OUI enrichment.

— Single named env var with clear but limited scope description v26.10.0

10 **Corelight Zeek detections for php2shell and weak eMudhra keys** NEW 40

Two Zeek detections authored by Corelight are added: `php2shell-detect` for the php2shell vulnerabilities and `emudhra-weak-keys` for weak eMudhra certificate keys.

— Names the two detections but gives no further mechanism v26.10.0

└──▷ **BREAKING ON UPGRADE**

! The environment variables `ZEEK_SYNCHROPHASOR_PORTS`, `ZEEK_GENISYS_PORTS`, and `ZEEK_ENIP_PORTS` are renamed to `ZEEK_PORTS_SYNCHROPHASOR`, `ZEEK_PORTS_GENISYS`, and `ZEEK_PORTS_ENIP`; existing values are migrated on upgrade, but any external automation referencing the old names must be updated.

WAS THIS USEFUL? 👍 👎

Asimily

Feedback ⓘ announced · Oct 28, 2026 BLOG ↗

Asimily is the proactive cyber asset defense platform for IoT, OT, IT, and IoMT devices, with prioritized risk mitigation.

Asimily Edge now deploys as a Cisco IOx container on Catalyst switches for deep packet inspection of east-west traffic without a separate appliance.

NEW INTEGRATION Asimily Edge now runs as a containerized IOx application on Cisco Catalyst switches via Cisco IOx application hosting, replacing the need for a separate physical or virtual data collection appliance on those switch segments.

Deploy Asimily Edge to many switches at once from Cisco Catalyst Center, or install it on a single switch through the switch web interface or command line.

Asimily Edge IOx application defaults to 2,100 CPU units, 1,024 MB memory, and 25,000 MB disk, with resource allocation adjustable from the switch command line for a specific model.

On enrollment, the IOx application registers with the customer's tenant and appears on the Edge dashboard, typically within minutes.

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Sophos Central Firewall

Feedback ⓘ

2026.44Released on October 29, 2026 NOTES ↗ RANK

Sophos Central Firewall is a cloud-managed network security solution that provides firewall protection, intrusion prevention, and threat monitoring for organizations, administered through the Sophos Central platform.

> Sophos Central Firewall adds asynchronous processing for SD-WAN connection group operations, along with a Force Delete option for stuck groups.

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

01 Asynchronous SD-WAN connection group management NEW

65

Create, edit, and delete operations on SD-WAN connection groups are now processed asynchronously in the background, with progress visible on the SD-WAN Connection and Task Queue pages. A new Force Delete option lets admins remove connection groups that cannot be deleted through the normal delete flow.

Monitor the status of a large SD-WAN group deletion in progress rather than waiting for a synchronous response.

> 📍 In the Sophos Central Firewall console, go to SD-WAN Connection page › select the connection group › choose Delete, then track progress on the Task Queue page.

Remove a stuck SD-WAN connection group that cannot be deleted through the normal delete flow.

> 📍 In the Sophos Central Firewall console, go to SD-WAN Connection page › select the stuck group › choose Force Delete.

— Names exact UI pages and Force Delete option, but no API/CLI detail.

2026.44Released on October 29, 2026

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ

IAM role fact check says Unknown when it cannot test the trust: NOTES ↗

Plerion is a cloud security platform that detects misconfigurations, vulnerabilities, and threats across cloud environments and provides remediation guidance.

Plerion adds SAST scanning for GitLab and Azure DevOps integrations and a new AWS SHIP lens on the Well-Architected page.

NEW DETECTION Enable static application security testing for GitLab and Azure DevOps integrations to scan Java, Kotlin, TypeScript, Python, Go and other source files on default branch and merge or pull request scans, reporting injection flaws and weak cryptography as findings and inline comments; off by default.

NEW COMPLIANCE FRAMEWORK Plerion's Well-Architected page now includes the AWS Security Health Improvement Program (SHIP) lens, grouping AWS Well-Architected Security best practices under SHIP's 11 security use cases with answers populated from Plerion findings.

UPDATED VULNERABILITY CHECK IAM role publicly accessible fact check (PLERION-AWS-553) — Unknown state for untestable trust relationships

RETIRED VULNERABILITY CHECK Fourteen publicly-accessible fact checks retired — decisions now derived from configuration alone

UPDATED VULNERABILITY CHECK EBS snapshot publicly accessible fact check (PLERION-AWS-244) — batched queries and longer retries

UPDATED VULNERABILITY CHECK EC2 instance publicly accessible fact check (PLERION-AWS-194) — active HTTP probe through ALB listeners

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.47.0 NOTES ↗ CODE ↗ RANK

This repo hosts Prowler OSS (CLI and Local Server). Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment for Humans and AI Agents.

> Prowler 5.47.0 adds new Azure compliance checks for Application Gateway WAF/HTTP2 and subscription owner counts, annotates MCP Server tools with read/write intent hints, redesigns the Issues list around exploitability, unifies the add-provider wizard across all supported clouds, and improves AWS security group exposure detection to catch combined-CIDR rules.

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

01 **New Azure checks for App Gateway and subscription owners** NEW

78

Prowler adds `network_application_gateway_waf_enabled` and `network_application_gateway_http2_enabled` checks for Azure Application Gateways and their WAF policies, plus `iam_subscription_owners_between_2_and_3` which verifies each subscription has between two and three principals with the Owner role, including owners inherited from management groups. All three new checks are mapped to CIS 5.0 and 6.0.

— Exact check IDs and CIS mapping named, mechanism explained. 5.47.0

02 **Combined-CIDR internet exposure detection for AWS security groups**

IMPROVED

78

Prowler's AWS security group internet exposure checks now treat CIDRs in a single ingress rule that together cover `0.0.0.0/0` or `::/0` (such as `0.0.0.0/1` plus `128.0.0.0/1`) as open to the internet.

— Exact CIDR mechanism and worked example given. 5.47.0

03 **Read-only/destructive hints in MCP Server tools** NEW

75

Prowler's MCP Server now annotates every tool with `readOnlyHint` and `destructiveHint` so MCP clients can distinguish read-only listing calls from writes; mute rule and mutelist writes are marked destructive.

— Exact annotation field names and behavior given. 5.47.0

04 **Unified add-provider wizard for all cloud providers** IMPROVED

65

Prowler's add-provider wizard now connects every provider — AWS, Azure, GCP, Microsoft 365, GitHub, Alibaba Cloud, and Cloudflare — in one step, with teammate invitations available for all providers and Lighthouse AI assistance beside the wizard.

— Lists every supported provider but no API or flag names. 5.47.0

THINNER COVERAGE BELOW

05 **Exploitability issues list redesign** IMPROVED 55

Prowler's Issues list now leads with the verdict, shows the attack path under each title, and adds a Why column with five evidence legs; the menu entry is renamed to Exploitability issues.

— UI changes described but no named config or API surface. 5.47.0

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - October 29, 2026 BLOG ↗

RANK

Sysdig Secure is a cloud security platform that provides runtime threat detection, vulnerability management, and compliance monitoring for containers, Kubernetes, and cloud workloads.

> Sysdig Secure adds preview automations that group Vulnerability Findings notifications by resource or image instead of sending one per finding.

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

01 **Grouped Vulnerability Findings automations** NEW 60

Sysdig Secure now supports Vulnerability Findings automations that group findings by resource or image, sending a single Slack or email notification — or creating a single Jira ticket — per resource or image rather than per finding. This feature is in Preview.

.
.
.

— Describes behavior and scope but no config key or API to act on directly.

Secure SaaS - October 29, 2026

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

cilium-cli

Feedback ⓘ v0.20.2 NOTES ↗ RANK

CLI to install, manage & troubleshoot Kubernetes clusters running Cilium

> cilium-cli v0.20.2 improves connectivity testing through restrictive proxies and namespace labeling, while removing deprecated BGP CRD support.

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

01 **Removal of deprecated v2alpha1 BGP CRDs** BREAKING 55

The deprecated `v2alpha1` version of the BGP CRDs is removed; clusters using `v2alpha1` BGP CRDs must migrate before upgrading.

— Names exact CRD version requiring migration before upgrade v0.20.2

02 **WebSocket fallback for SPDY-rejecting proxies** IMPROVED 45

cilium-cli now falls back to WebSocket when a proxy rejects the SPDY upgrade, allowing connectivity tests to work through proxies that block SPDY.

— Explains mechanism but no flag or command given v0.20.2

03 **Labels applied to all new test namespaces** IMPROVED 30

cilium-cli now applies labels to all new test namespaces, making it easier to identify and filter test resources.

— No label keys or filter syntax specified v0.20.2

└──▷ **BREAKING ON UPGRADE**

! The deprecated `v2alpha1` version of the BGP CRDs is removed; clusters using `v2alpha1` BGP CRDs must migrate before upgrading.

WAS THIS USEFUL? 👍 👎

security-profiles-operator

Feedback ⓘ v1.1.2

NOTES ↗ RANK

The Kubernetes Security Profiles Operator

> security-profiles-operator v1.1.2 adds a `spoc sign` subcommand with SLSA Build L3 provenance for its bundle, catalog, and security profiles, plus a new SPOD option to reject Permissive-mode SELinux profiles, alongside several breaking changes to metrics auth, Helm validation, profile policy, and resource naming.

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

01 Reject Permissive-mode SELinux profiles NEW

75

SelinuxProfiles that run in Permissive mode — which bypass the SELinux denylist — can now be rejected cluster-wide by setting `spec.selinux.options.allowPermissiveProfiles: false` in the SPOD.

Block Permissive-mode SELinux profiles from bypassing the denylist across the cluster.

yaml

```yaml
spec:
  selinux:
    options:
      allowPermissiveProfiles: false
```

— Exact config key and effect given, with runnable YAML example. v1.1.2

02 Metrics access requires spo-metrics-client ServiceAccount

75

BREAKING

Metrics access moves from the `default` ServiceAccount to the new `spo-metrics-client` ServiceAccount; scrapers using the `metrics-token` Secret must switch to `spo-metrics-client-token` or obtain a token via `kubectl create token spo-metrics-client`.

— Names exact ServiceAccounts, Secret, and migration command. v1.1.2

03 Sign OCI artifacts with spoc sign NEW

65

security-profiles-operator ships SLSA Build L3 provenance for the bundle, catalog, and security profiles, and adds a new `spoc sign` subcommand to sign OCI artifacts (e.g. `spoc sign <oci-artifact-ref>`) after recording them, for supply-chain verification.

Sign a security profile OCI artifact after recording it, for supply-chain verification.

```
$ spoc sign <oci-artifact-ref>
```

— Names the subcommand and provenance level but not verification mechanism. v1.1.2

THINNER COVERAGE BELOW

04 CIL keyword restrictions in SelinuxProfiles

BREAKING

55

SelinuxProfiles that use the CIL keywords `all`, `and`, `not`, `or`, `range`, or `xor` as permissions are now rejected.

— Lists exact rejected keywords but not remediation steps. v1.1.2

05 SPOD resources must be named spod

BREAKING

37

New SecurityProfilesOperatorDaemons must be named `spod`; other names are now rejected.

— Clear constraint but no further mechanism or migration detail. v1.1.2

06 Helm chart rejects unknown values

BREAKING

28

The Helm chart now rejects unknown values, so previously-ignored typos or stale keys in Helm chart configuration will now cause install/upgrade failures.

— No specific keys or migration steps given. v1.1.2

└──▷ BREAKING ON UPGRADE

- ! Metrics access moves from the `default` ServiceAccount to the new `spo-metrics-client` ServiceAccount; scrapers using the `metrics-token` Secret must switch to `spo-metrics-client-token` or use `kubectl create token spo-metrics-client`.
- ! The Helm chart now rejects unknown values.
- ! SelinuxProfiles that use the CIL keywords `all`, `and`, `not`, `or`, `range`, or `xor` as permissions are rejected.
- ! New SecurityProfilesOperatorDaemons must be named `spod`.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Anti-Virus for macOS

Feedback ⓘ

Version 2026.3.0 NOTES ↗ RANK

Sophos Anti-Virus for macOS is a security product that scans Mac systems for malware, viruses, and other threats, providing real-time protection and remediation.

> Sophos Anti-Virus for macOS gained AI Defense integration via Sophos Fusion and a new Web Profiles system with Gen AI categorization that replaces classic web control.

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

01 **Web Profiles with Gen AI categorization** NEW 43

Sophos Anti-Virus for macOS now supports Web Profiles with Gen AI categorization, replacing the previous classic web control system.

— States replacement and category but no configuration path given Version 2026.3.0

02 **AI Defense integration via Sophos Fusion** NEW 40

Sophos Anti-Virus for macOS now supports AI Defense through Sophos Fusion, providing visibility and control over AI usage.

— Names integration but no setup steps or config details Version 2026.3.0

WAS THIS USEFUL? 👍 👎

SentinelOne

Feedback ⓘ announced · Oct 28, 2026 PRESS ↗

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

SentinelOne adds Singularity ASOC, a new autonomous SOC suite with agentic investigation, hyperautomation, AI SIEM, IOC management, and a credits-based consumption model.

NEW INTEGRATION Singularity ASOC bundles Agentic Investigation, Hyperautomation, AI Data Pipelines, AI SIEM, and threat intelligence and IOC management into one suite on the Singularity Platform.

Singularity Credits now govern AI-powered capability consumption across the platform, with administrators setting guardrails to keep automated actions within organisation-defined limits.

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Agentless Vulnerability & Threat Detection version 1.651.10191318 (2026-10-19)

BLOG ↗ RANK

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

> Trend Vision One's Agentless Vulnerability & Threat Detection update lets AWS customers keep agentless scanning traffic inside their own VPC instead of routing through public endpoints.

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

01 **VPC-scoped agentless scanning deployment for AWS**

NEW SUPPORTED PLATFORM IMPROVED 60

Agentless Vulnerability & Threat Detection can now be deployed so scanning traffic stays within your own VPC rather than routing through public endpoints. In the console, go to Cloud Security → Cloud Accounts → AWS and redeploy using the CloudFormation template, selecting your target VPC and subnet when prompted.

When you need to restrict agentless scanning traffic to your own VPC instead of routing it through public endpoints, deploy the CloudFormation stack into your VPC via Cloud Security → Cloud Accounts → AWS.

> 📍 In the console, go to Cloud Security → Cloud Accounts → AWS, then redeploy using the CloudFormation template, selecting your target VPC and subnet when prompted.

— Only a usage example describes it; no mechanism beyond redeploy steps.

Agentless Vulnerability & Threat Detect…

> Agentless Vulnerability & Threat Detection version 1.651.10191318 (2026-10-19)
>
> BLOG ↗ STABLE
>
> Adds `SentryBucketPublicAccessBlockOmit` CloudFormation parameter (default `false`) to support SCP-restricted AWS environments in the AI Security stack.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Microsoft Graph Logs connector is available in Beta (2026-10-30) BLOG ↗ RANK

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

> Databricks shipped GA dashboard relationships for multi-fact AI/BI data models, a round of Genie Agents UX and performance improvements, Beta support for running multiple real-time flows per serverless pipeline alongside adaptive source-stage repartitioning for Kafka/Kinesis, and new `ug usage` CLI commands for spend-limit visibility.

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

01 **Spend limit visibility in ug CLI** NEW 80

View up to five personal and shared spend limits sorted by highest percentage used with `ug usage`, or see all applicable limits with `ug usage --all`.

See every personal and shared spend limit in your account to audit which budgets are close to exhaustion.

```
$ ug usage --all
```

— Exact runnable commands with concrete limit of five

View personal and shared spend limits w…

02 **Source stage adaptive repartitioning for streaming sources**

NEW 73

Source stage adaptive repartitioning automatically adjusts the number of tasks reading from Kafka or Kinesis sources to match traffic, removing the need to manually set `maxPartitions` or `shardsPerTask`.

— Names config keys it replaces, clear before/after

Multiple real-time flows per serverless…

03 **Genie Agents creation and file handling improvements** IMPROVED

65

Genie Agents gains a guided creation experience that walks through building an effective agent, retrieves files faster when answering questions (Beta), hides content search details such as the number of indexed files from users with `CAN VIEW` or `CAN`

`RUN` permission on a Genie Agent (Beta), and now opens links to volume files in an embedded agent directly instead of redirecting to the workspace.

— Names permissions and behaviors but no commands to invoke

Genie Agents enhancements and bug fixes…

THINNER COVERAGE BELOW

04 **Multiple real-time flows per serverless pipeline** NEW

55

You can now run multiple real-time mode flows within a single serverless pipeline without the flows competing for task slots; this is available on the preview channel in Beta.

— Explains behavior change but no setup steps given

Multiple real-time flows per serverless…

05 **Dashboard relationships in AI/BI** IMPROVED

40

Dashboard relationships are now generally available, letting you build complex multi-fact data models scoped to an AI/BI dashboard.

— No mechanism or UI path given beyond GA status Dashboard relationships are now general…

NEW INTEGRATION Microsoft Graph Logs connector (Beta) in Lakeflow Connect

NEW INTEGRATION Xero connector (Beta) in Lakeflow Connect

└──▷ ALSO FROM THESE RELEASES

> Microsoft Graph Logs connector is available in Beta (2026-10-30) BLOG ↗
>
> LATEST STABLE

> Xero connector is available in Beta (2026-10-29) BLOG ↗ STABLE

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Manage Features for Your Organization (Manage) (2026-10-19) BLOG ↗ RANK

Sumo Logic is a cloud-native machine data analytics platform that collects, searches, and analyzes logs and metrics for security monitoring, troubleshooting, and operational insights.

> Sumo Logic introduced a self-service Manage Features capability, letting paid customers toggle organization-level features on or off directly from the console.

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

01 Self-service Manage Features in Organizations UI NEW

50

Paid customers can now opt in or out of specific Sumo Logic features for their organization directly from the Organizations UI, without needing to raise a support request. In the console, go to Organizations UI > Manage Features, then toggle the desired feature on or off.

Enable or disable a Sumo Logic feature for your organization without filing a support ticket.

> 📍 In the Sumo Logic console, go to Organizations UI › Manage Features, then toggle the desired feature on or off for your organization.

– Names exact UI path but lacks list of affected features.

Manage Features for Your Organization (…

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl Stream

Feedback ⓘ changelog · Oct 29, 2026 BLOG ↗

Cribl Stream is a vendor-neutral observability pipeline that routes, filters, transforms, and reduces machine data from logs and metrics before sending it to analytics or storage destinations.

Cribl Stream adds a Microsoft 365 Copilot Source and a Microsoft Fabric Real-Time Intelligence destination for telemetry collection and routing.

NEW INTEGRATION Cribl Stream adds a Microsoft Copilot Source that collects Microsoft 365 Copilot interaction data — including user prompts and Copilot responses — from the Microsoft 365 Copilot Interaction Export API, running scheduled collection jobs per user with pagination support and built-in progress tracking so later runs focus on newer interactions.

NEW INTEGRATION Cribl Stream adds a dedicated Microsoft Fabric Real-Time Intelligence destination, enabling Stream-managed data to be sent into Fabric real-time analytics workflows for event-driven scenarios, streaming data, and logs.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

defectdojo

Feedback ⓘ 3.4.304 NOTES ↗ RANK

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo's generic parser now supports Package URL (purl) identifiers for dependency findings, enabling direct links between findings and specific packages.

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

01 **Package URL support in generic parser** NEW 55

The generic parser now accepts a Package URL (`purl`) for a finding's dependency, letting you link findings directly to a specific package.

— Names the purl field and parser but no further mechanism or limits. 3.4.304

WAS THIS USEFUL? 👍 👎

oxo

Feedback ⓘ v2.21.0 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> oxo adds support for scanning Autonomous System Numbers as a new asset type, alongside a minor protobuf enhancement for TestFlight assets.

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

01 **Scan ASN assets via new subcommand** NEW 75

oxo can now scan Autonomous System Numbers directly using the new `oxo scan run asn` subcommand and the `asn` asset type, allowing recon scans to enumerate infrastructure associated with an ASN.

Run a recon scan against an Autonomous System Number to enumerate its associated infrastructure.

```
$ oxo scan run --install --agent agent/ostorlab/nmap asn AS15169
```

— Names exact subcommand and asset type with runnable example. v2.21.0

THINNER COVERAGE BELOW

02 **App name field added to TestFlight asset** IMPROVED 30

oxo's TestFlight asset now includes an app name field in its protobuf definition.

— Names the field but no usage detail given. v2.21.0

WAS THIS USEFUL? 👍 👎

◆ Patch Management

Automox

Feedback ⓘ

COMING SOON: Manage Exclusion Windows Directly by Device (2026-10-29) BLOG ↗

RANK

Automox is a check out our full list of release notes, and stay up to date with what

> Automox is moving exclusion window management from device groups to individual devices, adding a device-level UI workflow while deprecating the group-scoped scheduled window API endpoints.

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

01 **Scheduled window API endpoints deprecated for device-level model**

DEPRECATED 75

As exclusion management moves to the device level, Automox is updating or deprecating these public API endpoints: Create a New Scheduled Window , Get a Scheduled Window , Update a Scheduled Window , Delete a Scheduled Window , List Scheduled Windows (search) , 'Check Whether a Scheduled Window is Currently Active', Check Scheduled Window Status for Server Groups , List Next Scheduled Window Dates by Group , and List Next Scheduled Window Dates by Device . New per-device exclusion window API endpoints will replace these group-scoped endpoints for programmatic management.

— Lists all affected endpoints by name but no exact paths or methods.

COMING SOON: Manage Exclusion Windows D…

02 **Device-level exclusion window management**

BREAKING 65

Admins can now assign and manage exclusion windows directly on individual devices via a new Edit Devices action in the Automox console, replacing the group-association workflow. Device tables and Device Details pages now show which exclusion windows apply to each specific device. As a behavior change, moving a device between groups no longer adds or removes its exclusion window, so any group-based scripts or onboarding flows that rely on group placement to apply or remove exclusion windows must be updated to manage exclusion windows as a separate, independent step.

Verify exclusion window assignments for a device after migration by reviewing the Device Details page in the Automox console.

📍 In the Automox console, go to Devices › select a device › Device Details page, and review the exclusion windows listed for that device to confirm expected

coverage post-migration.

— Names UI action and pages but not full navigation steps or config.

COMING SOON: Manage Exclusion Windows D…

└──▷ BREAKING ON UPGRADE

! The following public API endpoints are being updated or deprecated as exclusion management moves to the device level: Create a New Scheduled Window, Get a Scheduled Window, Update a Scheduled Window, Delete a Scheduled Window, List Scheduled Windows (search), 'Check Whether a Scheduled Window is Currently Active', Check Scheduled Window Status for Server Groups, List Next Scheduled Window Dates by Group, and List Next Scheduled Window Dates by Device.

! Moving a device between groups will no longer add or remove its exclusion window — any group-based scripts or onboarding flows that rely on group placement to apply or remove exclusion windows must be updated to manage exclusion windows as a separate, independent step.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ Iru Agent for Windows 1.25.7 NOTES ↗ RANK

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

> Iru Agent for Windows 1.25.7 strengthens the integrity of its download pipeline and improves app inventory visibility on Windows devices.

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

01 **Checksum validation for agent downloads** IMPROVED 40

Iru Agent for Windows now validates the checksum of every app and update it downloads, installing only files that match what Iru sent.

— Explains mechanism but no config surface to act on Iru Agent for Windows 1.25.7

02 **Expanded install-path reporting in app inventory** IMPROVED 28

App inventory for Windows computers now reports the install path for more applications.

— States the change but no scope or specifics Iru Agent for Windows 1.25.7

WAS THIS USEFUL? 👍 👎

Microsoft Intune

Feedback ⓘ Week of October 26, 2026

NOTES ↗

Microsoft Intune is a cloud-based unified endpoint management solution built on Zero Trust. Empower IT to manage, secure and govern devices, apps and endpoints.

Microsoft Intune adds bulk eSIM activation for up to 100 corporate-owned Android Enterprise devices, with eSIM data plan preservation during bulk wipes.

NEW SUPPORTED PLATFORM Activate eSIMs on up to 100 corporate-owned Android Enterprise devices at once via `Devices > All devices > Bulk device actions`, supplying a carrier activation server URL, on Android 15 or later (COBO, COSU, and COPE form factors).

When bulk wiping supported devices, Intune now preserves eSIM data plans by default, with an opt-in option to also remove the data plans during the wipe.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ runZero 5.1.261029.0 (2026-10-29) BLOG ↗

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

runZero adds `has_owner` search to the vulnerability inventory and avoids probing Allen-Bradley/Rockwell 1785-ENET devices on identification.

Search the vulnerability inventory by asset owner using the new `has_owner` search term.

runZero now avoids sending traffic to Allen-Bradley/Rockwell 1785-ENET devices as soon as they are identified during scanning.

NEW VULNERABILITY CHECK eMudhra compromised certificate vulnerability check

UPDATED DETECTION Fingerprint improvements for Beckhoff Automation

UPDATED DETECTION Fingerprint improvements for Siemens and Rockwell Automation

WAS THIS USEFUL? 👍 👎

Axonius

Feedback ⓘ Axonius Release Notes 10.0.1 (2026-10-26) BLOG ↗

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

Axonius 10.0.1 adds AI Workforce auto-pause, a new Adapter Noise Filters Worker, AVS factor querying, Security Finding Risk Score v2, and scoped Vulnerability Repository access.

AI Workforce scheduled workers now pause automatically after 14 consecutive runs with the same result and no action taken; a notification card in AI Workforce > Notifications explains the pause and offers Resume.

Adapter Noise Filters Worker is a new weekly AI Workforce worker in the IT Asset Management department that finds uncorrelated devices and identifies which adapters bring them in, pointing to Advanced Settings for remediation.

Query AVS factors as structured fields on the Aggregated Security Findings and Vulnerability Repository pages using the Query Wizard, with autocomplete for each factor's Impact and value, plus support for custom values.

Security Finding Risk Score v2 now weights Axonius Vulnerability Score at 60% (up from 50%) and Asset Criticality at 35% (up from 25%); activating v2 deactivates v1 without deleting it, and scores recalculate immediately.

Scoped users and admins in a data scope can now open the Vulnerability Repository and use its table, filters, queries, charts, and export.

+2 more

NEW INTEGRATION CMDB Reconciliation Workspace adds six new ITSM adapter sources

└─▷ **BREAKING ON UPGRADE**

! Activating Security Finding Risk Score v2 deactivates v1; anything built on the previous scores (queries, dashboards, enforcement sets) may show different risk levels immediately after activation.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

LoginRadius

Feedback ⓘ changelog · Oct 30, 2026 CHANGELOG ↗

RANK

LoginRadius: Privacy-compliant identity and access management platform. Offer Authentication and Authorization for your web and mobile apps in just a few clicks, with: SSO, MFA, social logins and more.

> LoginRadius enhanced its Admin Console with direct Phone ID management for user records and tightened permission restrictions for view-only users.

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

01 **Phone ID management in Admin Console** NEW 65

Admins can add, update, or remove a user's Phone ID directly from `User details > Edit User Info`. The field only appears when phone login is enabled, and changing the Phone ID marks it as unverified.

— Names exact UI path and behavior, no API/flag given changelog-20261030-5f560e78

THINNER COVERAGE BELOW

02 **View-only restriction in Organizations and Custom IdP**

IMPROVED 45

View-only users are now restricted from creating or editing items under Organizations and `Custom IdP > Settings`.

— Names settings areas but lacks mechanism detail changelog-20261030-5f560e78

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ changelog · Oct 30, 2026 CHANGELOG ↗ RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud's biggest window moves were a new SIEM integration for Directory Insights, expanded Apple MDM Device Lock with automatic PIN escrow, a Scrut Automation integration for continuous audit evidence, and a broad Agentic IAM release adding non-human identity lifecycle management, continuous audit trails, real-time risk detection, Conditional-Access device trust, workflows automation, and an upcoming PAM module.

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

01 **Device Lock enhancements for Apple MDM** IMPROVED 80

The Device Lock MDM command for macOS now automatically generates and escrows the lock PIN, retrievable from Admin Portal under `Highlights > Device Lock > View PIN`. Admins can show a custom message and phone number on the Lock Screen for Mac, iPhone, and iPad, and track lock status from the Devices list via the new `macOS Lock Status` column.

— Exact admin navigation and column name given. changelog-20261030-bcfbc4d0

02 **Scrut Automation integration for compliance evidence** NEW 65

JumpCloud now integrates with Scrut Automation so that active identity and device settings across Windows, Mac, and Linux endpoints are automatically collected and mapped to compliance framework requirements, removing manual screenshot-gathering during audits. AI agent ownership and access tracking also feeds into Scrut to produce auditor-ready evidence that agents follow required safety rules, across 70+ compliance frameworks.

— Names endpoints covered and framework count but no setup steps.

launch-20261029-428781ad

THINNER COVERAGE BELOW

03 **Continuous audit trails for human and non-human actions** NEW 45

Continuous audit trails now link every human and non-human action to a verified identity, supporting SOC 2, ISO 27001, HIPAA, and NIS2 audits.

— Names compliance frameworks but no access path. launch-20261028-a7780d5e

04 **Device trust via Conditional Access Policies** NEW 45

Device trust now blocks unmanaged or unhealthy devices from accessing company data using Conditional Access Policies (CAP).

— Names CAP mechanism but no configuration detail. launch-20261028-a7780d5e

05 **Agentic IAM for non-human identity lifecycle** NEW 40

Agentic IAM provides full visibility and lifecycle tracking across AI agents and non-human identities, automatically updating agent access when a user changes roles or leaves.

— Describes behaviour but no config, UI path, or limits. launch-20261028-a7780d5e

06 **Privileged access management for admins and agents** NEW 40

Privileged access management (PAM) will deploy isolated identities and time-boxed credentials for admins and agents, coming December 2026.

— Future feature with mechanism noted but not yet available. launch-20261028-a7780d5e

07 **Real-time risk detection for identities and agents** NEW 30

Risk detection flags risky sign-ins, non-human behavior, and unapproved access changes in real time.

— Generic description, no named surface or threshold. launch-20261028-a7780d5e

08 **Workflows automation for lifecycle and remediation** NEW 30

Workflows automation now automates lifecycle processes, offboarding, and policy remediation across all identities and devices.

— Broad prose claim without named workflow or trigger. launch-20261028-a7780d5e

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ changelog · Oct 30, 2026 CHANGELOG ↗ RANK

All the latest releases from the 1Password team

> 1Password CLI 2.43.0-beta.02 adds a command to list secrets stored in Environments, alongside a small documentation fix for file attachment reads.

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

01 **List secrets from Environments via CLI** NEW 75

The new `op environment secret list` command, added in CLI 2.43.0-beta.02, reads secrets from Environments, letting users audit stored values or pipe them into scripts.

List all secrets stored in an Environment to audit or pipe values into a script.

```
$ op environment secret list
```

— Names exact command and version; runnable example given changelog-20261030-a320ab24

THINNER COVERAGE BELOW

02 **Help text example for reading file attachments** IMPROVED 35

The `op read` command's help text now includes an example demonstrating how to read a file attachment.

— Minor doc update, no mechanism or example given changelog-20261030-a320ab24

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ Role mining suggestions (2026-10-29)

BLOG ↗ RANK

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

> ConductorOne's single release this window adds role mining suggestions that proactively identify candidate groups for role creation and let analysts automate membership enrollment for new hires.

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

01 **Role mining suggestions with automated membership enrollment**

NEW 73

ConductorOne surfaces role mining suggestions showing each group's size, shared entitlements, and a confidence score based on attributes such as department, job title, or manager. Selecting a suggestion opens an analysis page with the group's filters pre-filled, letting reviewers filter by user attribute or by access held and limit scope to specific apps; reviewers can then turn on membership automation so users matching the group's filters — including new hires — receive the role's access automatically. Any earlier role mining analysis can be resumed from the start page.

When onboarding a new department, use a role mining suggestion to quickly create a role that auto-enrolls future hires.

> 📍 1. Go to Role Mining in the C1 console. 2. Review the suggestions list — each entry shows group size, shared entitlements, and confidence score. 3. Pick a suggestion and select `Open analysis` — the group's filters are pre-filled. 4. Filter by user attribute (e.g. department = `Engineering`) or by specific app access to refine membership. 5. Create the role, then enable membership automation so new hires matching the filters receive access automatically.

— Clear mechanism and UI steps given, but no API/config names or limits.

Role mining suggestions (2026-10-29)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

HTTP Message Signatures Added to Tenant Network ACLs (2026-10-29) BLOG ↗

RANK

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

> Auth0 brings RFC 9421 HTTP Message Signatures to Tenant Network ACLs, giving enterprise tenants cryptographic verification of request origin at the network boundary.

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

01 HTTP Message Signatures in Tenant Network ACLs NEW

88

Auth0 Tenant Network ACLs now support RFC 9421 HTTP Message Signatures via the Management API, letting enterprise tenants enforce HMAC-SHA256 request verification at the network boundary to block direct tenant bypass. Incoming requests can be validated to originate only from authorized reverse proxies or gateways, with unverified direct traffic dropped. Signing keys can be rotated with zero downtime by configuring multiple active signing keys per rule, and key-rotation activity and verification events can be observed via `acls_summary` Tenant Log events.

— Names RFC, API, HMAC method and log event key but no exact endpoint path.

HTTP Message Signatures Added to Tenant…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ Passkeys as a second factor (2026-10-29) BLOG ↗

RANK

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

> Clerk now supports passkeys as a second factor, allowing users to complete multi-factor authentication with a passkey after an initial sign-in method.

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

01 **Passkeys as a second factor** NEW 40

Clerk now supports passkeys as a second factor for multi-factor authentication, letting users verify with a passkey after signing in with another method.

— States what changed but no setup steps or API/config named

Passkeys as a second factor (2026-10-29)

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Version 2026.44.0 (2026-10-28) BLOG ↗

Okta identity engine is a secure, scalable, and highly available authentication and user management for any app.

Okta Identity Engine adds MCP server OAuth client and Cross App Access support, High Risk Proxy detection, passkey AAGUID-only group matching, and new Privileged Access System Log events.

Match passkeys to authenticator groups on AAGUID alone with the new self-service EA setting in Preview, without requiring a full attestation match.

NEW INTEGRATION Connect MCP servers to downstream resources as OAuth 2.0 clients, now GA in Preview.

NEW INTEGRATION Enable Cross App Access for MCP servers, now GA in Preview.

NEW INTEGRATION Enable Cross App Access with an external identity provider, now GA in Preview.

Actor claim support for OAuth 2.0 On-Behalf-Of Token Exchange is now GA in Production.

+4 more

NEW DETECTION High Risk Proxy detection

NEW SUPPORTED PLATFORM Anything-as-a-Source device import GA in Production

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ changelog · Oct 30, 2026

CHANGELOG ↗ RANK

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

> DigiCert CertCentral added API-level visibility into supported intermediate CA chains and automated key-type handling for certificate orders through its Europe Services API.

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

01 **Automatic key-type matching in Europe Services API** NEW

85

When ordering through the CertCentral Europe Services API (`https://certcentral.digicert.eu/services/v2`), if the optional `ca_cert_id` parameter is omitted on a product set to the DigiCert default ICA chain, the system automatically performs key-type matching between the ICA chain and the submitted CSR.

— Names endpoint, parameter and exact trigger condition changelog-20261030-373314c3

02 **ICA chain visibility via Product info API** NEW

65

The Product info API now exposes an `allowed_ca_certs` field listing the intermediate CA (ICA) chains supported for a given product. If the field is absent from the response, ICA selection is not enabled for that product.

— Names exact field but no sample request/response shown changelog-20261030-373314c3

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ changelog · Oct 30, 2026 CHANGELOG ↗ RANK

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

> Tailscale introduced a new route selection policy to replace Regional Routing, added per-machine posture assertion visibility in the admin console, and shipped a Terraform breaking change alongside a minor macOS UI fix.

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

01 **Terraform split_dns field now a set** BREAKING 70

The `split_dns` field in the `tailscale_dns_configuration` Terraform resource is now defined as a set instead of a list, which may require updates to existing configurations.

– Names exact field and resource but no migration example changelog-20261030-17a31fd1

02 **Route selection policy setting (MagicRoute)** NEW 65

A new Route selection setting in the admin console replaces the Regional Routing toggle, letting admins control how the tailnet chooses between redundant subnet routers, app connectors, exit nodes, and Tailscale Service hosts. Configure it via Settings > Route selection to set MagicRoute preferences.

When you have redundant subnet routers or exit nodes and want to control which path the tailnet prefers, configure Route selection in the admin console.

> 📍 In the admin console, go to Settings › Route selection and configure MagicRoute to choose between redundant subnet routers, app connectors, exit nodes, and Tailscale Service hosts.

– Names UI path and mechanism but lacks deeper config detail changelog-20261030-17a31fd1

03 **Per-machine posture assertion details** NEW 60

The postures page of the admin console now shows posture assertion details per machine, including each machine's value, pass/fail status, and a count of policy file rules that require a given posture.

– Describes what's shown and where, no command changelog-20261030-17a31fd1

THINNER COVERAGE BELOW

04 **Notched-display notification for hidden icon** IMPROVED

MacBooks with a notched display now show a notification window when the Tailscale menu bar icon is hidden behind the notch.

— Minor UI fix with no further mechanism detail `changelog-20261030-17a31fd1`

└──▷ **BREAKING ON UPGRADE**

! The `split_dns` field in the `tailscale_dns_configuration` Terraform resource is now defined as a set, not a list, which may require updates to existing configurations.

WAS THIS USEFUL?  

Slack Nebula

Feedback ⓘ v1.12.0 NOTES ↗ CODE ↗ RANK

A scalable overlay networking tool with a focus on performance, simplicity and security

> Nebula v1.12.0 adds Linux UDP/TUN offloads, CPU thread pinning, a local control API reachable via `nebula ctl`, an HTTP info socket for node identity, FIPS 140-3 build mode, and experimental split DNS, alongside a large set of breaking changes to config validation, SSH debug server behavior, and the embedder API.

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

01 Linux UDP/TUN offload support NEW

90

Nebula can offload packet processing to the Linux kernel via `tun.use_offloads` (enables TSO/USO) and `listen.udp_offloads` (enables UDP GSO/GRO) in config, reducing CPU overhead at high throughput; both are off by default but will become the default in a future release.

Opt into Linux kernel offloads to reduce CPU overhead on high-throughput Nebula nodes.

yaml

```yaml
tun:
  use_offloads: true

listen:
  udp_offloads: true
```

— Names exact config keys and offload types with a runnable example v1.12.0

02 CPU thread pinning for I/O routines NEW

90

When `routines` is set above 1 on Linux, each I/O routine's thread can be pinned to a CPU using `tun.pin_threads` (default `true`), `tun.pin_threads_key` (`pid` or `port`), and `tun.cpu_affinity`, reducing latency jitter on multi-core hosts.

Spread I/O routines across physical performance cores to reduce latency jitter on multi-core Linux hosts.

```yaml
routines: 4

tun:
  pin_threads: true
  pin_threads_key: pid
```

— Full mechanism and config keys with a working example v1.12.0

03 **FIPS 140-3 build mode** NEW 90

CGO-free FIPS 140-3 binaries can be built with `make fips140` or `make release-fips140`; these builds restrict operation to P256 certificates and AES, and `nebula-cert ca` and `keygen` default to `-curve P256`.

— Exact build commands and resulting cryptographic restrictions v1.12.0

04 **Embedder API breaking changes** BREAKING 90

Library embedders must update to new signatures: `config.C.RegisterReloadCallback` callbacks now take (*config.C, *config.Reload); `ReloadConfig` and `ReloadConfigString` return a `config.ReloadResult` (previously `ReloadConfigString` returned an `error`, now available as the result's `Err`); `overlay.Device` no longer embeds `io.ReadWriteCloser`; `NewMultiQueueReader` and `SupportsMultiqueue` are replaced by `Queues`; `udp.Conn` gained `WriteBatch` and `IPv4Only`; `ListenOut` takes a flush callback; `ReloadConfig` takes a `*config.Reload`; `udp.NewListener` and `udp.NewGenericListener` take a `udp.Settings`; SSH debug commands moved to a new `commands` package, `sshd.RegisterCommand` and `sshd.Command` are removed, `sshd.NewSSHServer` takes a `*commands.Registry`, and `firewall.ProtoTCP`, `ProtoUDP`, `ProtoICMP`, `ProtoICMPv6` are now `iputil.IPProtocol*`.

— Lists every renamed/changed API surface with before/after v1.12.0

05 **Local control socket via `nebula ctl`** NEW 85

Every SSH debug command is now reachable through a new local control socket using `nebula ctl <command>`; access is restricted to the running user, root, or a group named in `local_api.control.group`, and the socket path can be set per-instance with `local_api.control.socket`.

Inspect the certificate and VPN identity of a peer from the local host without an SSH session.

```
$ nebula ctl print-cert
```

— Named command, config keys, and access model with an example v1.12.0

06 **Stricter config and startup validation** BREAKING 80

Firewall rules with a `group` or `groups` field that YAML reads as a float, date, map, or nested array are now rejected at startup and reload instead of being silently coerced to a string (previously `group: 1.10` matched group `1.1`), and a null `group` / `groups` is rejected when no `host`, other group, or `cidr` limits the rule. A `pki.blocklist` entry that isn't a hex SHA-256 fingerprint now causes Nebula to refuse to start or reload. On Linux, Nebula now refuses a `tun.dev` that already exists (including persistent devices from `ip tuntap add` or systemd-networkd). `tun.mtu` and any route `mtu` above 8937 is capped to 8937 with a warning, and `nebula` / `nebula-service` now refuse to start when given an argument that isn't a flag.

— Every validation rule and affected config key is named v1.12.0

07 **HTTP info socket for node identity** NEW 75

An opt-in info socket exposes node identity over HTTP: `GET /v1/identity` returns the local node's identity, `GET /v1/identity/<vpn addr>` returns another host's, and appending `/pem` to either path returns the certificate.

— Exact endpoints named but no runnable example given v1.12.0

08 **SSH debug server behavior changes** BREAKING 70

SSH debug server commands now exit with status 1 when they refuse, and report refusal as `{"ok":false,"reason":...}` when run with the newly-supported `-json` or `-pretty` flags (now accepted by every command). The `start-cpu-profile`, `stop-cpu-profile`, `save-heap-profile`, `save-mutex-profile`, and `mutex-profile-fraction` commands, along with `sshd.sandbox_dir`, have been removed.

— Names exact removed commands and flags but no migration example v1.12.0

09 **Experimental split DNS** NEW 65

A new experimental `dns` stanza lets a node override its host resolver, pointing `dns.nameservers` at queries for `dns.match_domains` or all lookups; it's off by default and supported only on macOS and Windows.

— Config keys named but no example and marked experimental v1.12.0

THINNER COVERAGE BELOW

10 **Outbound UDP send batching** IMPROVED 45

Nebula now batches outbound UDP sends on Linux, issuing up to 128 packets per `sendmmsg` call to improve throughput.

— Automatic internal change with no user-facing control v1.12.0

11 **Go 1.27.3 build requirement** BREAKING 30

Building Nebula or embedding it as a library now requires Go v1.27.3.

— Single bare requirement with no further detail v1.12.0

└──▷ **BREAKING ON UPGRADE**

! Building Nebula or embedding it as a library now requires Go v1.27.3.

! A firewall rule `group` that YAML reads as a float, a date, a map, or a nested array is now rejected on startup and reload instead of being silently converted to a string (e.g. `group: 1.10` used to match group `1.1`). A null `group` or `groups` is rejected when no `host`, other group, or `cidr` limits the rule.

! A `pki.blocklist` entry that is not a hex SHA-256 fingerprint now causes Nebula to refuse to start and fail a reload.

! On Linux, Nebula refuses a `tun.dev` that already exists (including a persistent device created with `ip tuntap add` or systemd-networkd); remove it and let Nebula create the device.

! `tun.mtu` and any route `mtu` above 8937 is capped to 8937 with a warning.

! SSH debug server commands now exit with status 1 when they refuse; with `-json` or `-pretty` they report refusal as `{"ok":false,"reason":...}`. Every command now accepts `-json` and `-pretty`.

! `nebula` and `nebula-service` refuse to start when given an argument that is not a flag (previously ignored).

! Embedders: `config.C.RegisterReloadCallback` callbacks now take (*config.C, *config.Reload); `ReloadConfig` and `ReloadConfigString` return a `config.ReloadResult`; `ReloadConfigString` used to return an `error`, now it is the result's `Err`.

! Embedders: `overlay.Device` no longer embeds `io.ReadWriteCloser`; `NewMultiQueueReader` and `SupportsMultiqueue` are replaced by `Queues`; `udp.Conn` gained `WriteBatch`, `IPv4Only`; `ListenOut` takes a flush callback; `ReloadConfig` takes a `*config.Reload`; `udp.NewListener` and `udp.NewGenericListener` take a `udp.Settings`.

! Embedders: SSH debug commands moved to the new `commands` package; `sshd.RegisterCommand` and `sshd.Command` are removed; `sshd.NewSSHServer` takes a `*commands.Registry`; `firewall.ProtoTCP`, `ProtoUDP`, `ProtoICMP`, and `ProtoICMPv6` are now `iputil.IPProtocol*`.

! The SSH debug server's `start-cpu-profile`, `stop-cpu-profile`, `save-heap-profile`, `save-mutex-profile`, and `mutex-profile-fraction` commands, and `sshd.sandbox_dir`, are removed.

WAS THIS USEFUL?  

Firezone

Feedback ⓘ Gateway 1.6.4 (2026-10-29) BLOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone shipped performance improvements to WireGuard tunnel throughput and route handling, plus a redesigned Settings screen for choosing between US and EU control planes in the Linux GUI.

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

01 US/EU control plane selector in Linux GUI IMPROVED

50

A redesigned Settings screen in the Firezone Linux GUI lets users switch between the US or EU control plane, useful for organizations that need data to stay in European infrastructure.

Select the EU control plane endpoint if your organization's data must stay in European infrastructure.

> 📍 In the Firezone Linux GUI, go to Settings and select `EU` under the control plane region selector.

— Names exact UI path and both region options Linux GUI 1.5.20 (2026-10-29)

02 Parallelized WireGuard encryption across worker threads

IMPROVED

45

Firezone Gateway and the Linux GUI now parallelize WireGuard encryption and decryption across worker threads, increasing tunnel throughput for established connections.

— Explains mechanism but no benchmark numbers or config to act on

Gateway 1.6.4 (2026-10-29) Linux GUI 1.5.20 (2026-10-29)

03 Route resolution caching for established flows IMPROVED

30

Packet-routing performance in the Linux GUI now benefits from caching of route resolution for established flows.

— Thin one-line description with no mechanism detail or numbers

Linux GUI 1.5.20 (2026-10-29)

WAS THIS USEFUL?

◆ Data Security & DLP

Immuta

Feedback ⓘ updated · Oct 30, 2026 NOTES ↗ RANK

Use Immuta to provision governed access to data.

> Immuta expands subscription policy enforcement beyond its native data integrations with a new custom securables capability for external assets.

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

01 **Custom securables for external assets** NEW 70

Immuta now supports custom securables, letting you register external assets — dashboards, files, APIs, and SaaS applications — in Immuta and control access to them with subscription policies. Your tooling then calls Immuta's APIs or webhooks to retrieve access decisions and enforce them in the target system.

— Explains mechanism and asset types but no exact endpoint names given snapshot-20261030

WAS THIS USEFUL?

Varonis

Feedback ⓘ

Strengthen data and AI governance with Varonis + Collier NOTES ↗ RANK

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

> Varonis now supports Azure Managed Identity authentication for connecting managed databases, removing the need for manually managed service-account credentials.

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

01 **Azure Managed Identity for managed databases** NEW 55

Varonis can connect to Azure managed databases using Azure Managed Identity authentication instead of manually managed service-account credentials. Supported databases include Azure SQL Server, PostgreSQL, and MySQL.

— Names mechanism and three supported databases, no setup steps shown

Securely connect Azure managed database…

NEW INTEGRATION Collibra data catalog integration

> Strengthen data and AI governance with Varonis + Collier NOTES ↗ LATEST
>
> STABLE
>
> Varonis adds a Collibra integration that enriches catalog assets with sensitive data classification, security context, ownership, lineage, and policy information.

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

Aembit

Feedback ⓘ

Introducing the AI Agent Kill Switch: Instant Single-User Revocation (2026-10-29)

BLOG ↗ 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 introduced an AI Agent Kill Switch that lets admins instantly revoke a single user's MCP client session without disrupting other agents or users, with support for both targeted and bulk revocation plus API automation.

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

01 AI Agent Kill Switch for MCP session revocation NEW

88

Admins can revoke a specific MCP user session by navigating to Administration > Identity Providers > Sessions tab and selecting Revoke session, scoped to one user and one MCP client without affecting other sessions. Sessions in this view can also be filtered by user email, MCP client, or both and bulk-revoked using Revoke filtered, covering all sessions for one user or one MCP client. Revocation can be automated in incident-response workflows via the Aembit Cloud API. Sessions appear for MCP clients that connect through a Client-to-Gateway Access Policy using an OIDC ID Token Credential Provider with refresh token support enabled, and admins need `Read/Write` permission on Identity Providers to revoke.

Immediately cut off a single user's compromised Claude Code session without locking them out of other agents or breaking the tool for teammates.

📍 In your Aembit tenant: 1. Go to Administration › Identity Providers and open the provider your users sign in through. 2. Open the Sessions tab. 3. Filter by user email and MCP client to find the target session. 4. Select Revoke session and confirm.

Bulk-revoke all sessions tied to a retired MCP client during incident response, without touching any other user or agent.

📍 In your Aembit tenant: 1. Go to Administration › Identity Providers, open the relevant provider, and open the Sessions tab. 2. Filter by the retired MCP client name. 3. Select Revoke filtered and confirm.

— Names exact UI path, permissions, and prerequisites; API lacks endpoint detail

Introducing the AI Agent Kill Switch: I…

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 2.0.0 (2026-10-30) BLOG ↗ RANK

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

> Runlayer 2.0.0 hardens client-IP sourcing and session isolation to close cross-session manipulation and IP-spoofing gaps, while 1.41–1.42 add Terraform-managed skills with risk-policy enforcement, richer redacted session timelines, a compact tool-discovery header, and new Outlook, Google Docs, and LLM-budget-alert capabilities.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 Trusted-proxy client IP enforcement

BREAKING

88

Runlayer now determines client IP from the address your own proxy appended to `X-Forwarded-For`, never an earlier caller-supplied entry, so rate limits, `ip_address` on audit events, and the `meta.request.ip` policy field reflect one address per office or fleet rather than per workstation. Self-hosted deployments with more than one appending proxy must set `PRE_AUTH_TRUSTED_PROXY_HOPS` (Terraform: `trusted_proxy_hops`) to the number of proxies, and `CLIENT_IP_HEADER` is now honoured only when it carries a single IP address.

Tell Runlayer how many proxies prepend to X-Forwarded-For in your self-hosted deployment so IP-range policies and rate limits see the correct client address.

```env
PRE_AUTH_TRUSTED_PROXY_HOPS=2
```

— Exact env var, Terraform key, and behavior change documented. 2.0.0 (2026-10-30)

02 ExtendDB AWS credential handling in Kubernetes operator

85

BREAKING

The operator no longer injects `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY`, or `AWS_CA_BUNDLE` on MCP pods; ExtendDB SigV4 now uses AWS profile `extenddb` via `EXTENDDB_AWS_PROFILE` (credentials under `/etc/extenddb-aws`, CA on the profile), and custom `enable_db` images must build the DynamoDB client from that profile when `EXTENDDB_AWS_PROFILE` is set. The Kubernetes operator also now rejects MCP servers that combine `enable_db` with an IRSA annotation (`eks.amazonaws.com/role-arn`) on the MCP ServiceAccount, setting `ExtendDBReady=False` / `IRSAConflict` and leaving ExtendDB AWS env off the pod.

— Names every env var, path, and operator status condition. 1.41.0 (2026-10-29) 2.0.0 (2026-10-30)

03 **Redacted session timeline preservation** IMPROVED 85

Redacted sessions and agent runs now preserve more data instead of hiding it: the API and the `get_session` MCP tool return the full event timeline (times, event types, tool names, model, tokens) with `***REDACTED***` in place of prompts, responses, tool input/output, and page URLs; the `list_sessions` and `search_session_events` MCP tools no longer include the first prompt or generated title in list and search rows; and `get_agent_run_trace` keeps its steps with placeholder text instead of hiding the timeline entirely.

— Names every affected MCP tool and exact redaction behavior. 1.41.0 (2026-10-29)

04 **Compact tool discovery control header** NEW 83

Send `X-Runlayer-Tool-Exposure: force-initial` from any client to keep the Runlayer Plugin on `search_tools` / `execute_tool` instead of listing every connector tool, controlling compact tool discovery in agent harnesses.

Keep a Runlayer Plugin in compact tool-discovery mode for an agent harness that loads all tool definitions up front, avoiding the full connector tool listing.

```
$ curl -H 'X-Runlayer-Tool-Exposure: force-initial'
https://<your-runlayer-host>/mcp
```

— Exact header and a runnable request given. 1.41.0 (2026-10-29)

05 **Skill risk policy enforcement for API and Terraform** NEW 80

Skill risk policy now runs on API writes: creating or updating a skill's files via the API or Terraform is rejected when the skill risk policy blocks the content, with nothing saved when a write is blocked; the new `runlayer_skill` Terraform provider resource manages skills and their complete file sets as code, with changes made in the app appearing as a plan diff and a policy-blocked file set failing the apply.

Define and version-control a skill and its files in Terraform so that app edits surface as plan diffs and a blocked risk policy fails the apply.

```hcl
resource "runlayer_skill" "data_lookup" {
  name      = "data-lookup"
  namespace = "finance"
  path      = "skills/data-lookup"

  files = {
    "skill.yaml" = file("skills/data-lookup/skill.yaml")
    "handler.py" = file("skills/data-lookup/handler.py")
  }
}
```

— Describes enforcement point but no endpoint named. Exact resource name with a runnable Terraform example.

1.41.0 (2026-10-29)

06 **LLM budget email alerts** NEW 70

Set email alerts for LLM budget thresholds under Settings → LLM Gateway → Usage rules → Email alerts by configuring percentages; alerts fire once per percentage per budget window to organization settings administrators.

Configure LLM budget email alerts so administrators are notified before spend hits the full budget ceiling.

> 📍 In Settings, go to LLM Gateway › Usage rules, open an existing usage rule that has a total budget set, then open Email alerts and add the percentages (for example 50, 80, 100) at which administrators should be emailed.

— Exact UI path and trigger behavior given. 1.42.0 (2026-10-29)

07 **Terraform secret fields marked sensitive** BREAKING 70

The Terraform provider now marks `manual_oauth_client_secret` and `broker_vendor_client_secret` on `runlayer_server` as sensitive, hiding them in plan and apply output; module outputs that reference them now require `sensitive = true` or the apply fails.

— Names exact fields and the required fix. 1.41.0 (2026-10-29)

08 **Cross-session manipulation blocked** BREAKING 65

A tool call carrying someone else's session ID — via hook payload or the `x-runlayer-session-id` header — can no longer block that session, inject scan results into it, add security events to it, or mark an operator Stop as having taken effect in that session's history; hook events naming another session are no longer recorded into it.

— Clear mechanism and scope but no reader action required. 2.0.0 (2026-10-30)

09 **Identity-forward header stripping** BREAKING 60

Runlayer now drops `X-Runlayer-User-Email`, `X-Runlayer-Identity-Token`, and other identity-forward headers when a client sends them, so an upstream MCP server only ever sees the ones Runlayer injects; other custom headers such as a correlation ID still pass through.

— Names headers dropped but no remediation step given. 1.41.0 (2026-10-29)

THINNER COVERAGE BELOW

10 **model_version field in ToolGuard score API** NEW 55

ToolGuard score API responses now include a `model_version` field identifying the model version expected to score the request; include it when reporting scan results to Runlayer.

— Names the field but no endpoint path given. 1.42.0 (2026-10-29)

11 **Connector capability expansions (Outlook, Google Docs)** NEW 55

The Microsoft Outlook connector now supports creating mail folders, moving folders, and moving messages between folders, and the Google Docs connector now supports copying existing templates with native formatting preserved, inspecting document structure and styles, and attaching comments to specific text ranges; redeploy existing Outlook and Google Docs connectors to get the new tools.

— Lists all three new tools and redeploy requirement. Lists three new capabilities and redeploy requirement.

1.42.0 (2026-10-29) 1.41.0 (2026-10-29)

12 **Scheduled agent run results via iMessage and Slack** NEW 45

Scheduled agent runs set up from an iMessage thread now post each run's result back via iMessage, and those set up from a Slack thread now post results back in that thread.

— Describes behavior but no config surface. 1.42.0 (2026-10-29)

└──▷ **BREAKING ON UPGRADE**

! The operator no longer injects `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY`, or `AWS_CA_BUNDLE` on MCP pods; ExtendDB SigV4 now uses AWS profile `extenddb` via `EXTENDDB_AWS_PROFILE` (credentials under `/etc/extenddb-aws`, CA on the profile). Custom `enable_db` images must build the DynamoDB client from that profile when `EXTENDDB_AWS_PROFILE` is set.

! Client IP can no longer be chosen by the caller: rate limits, `ip_address` on audit events, and the `meta.request.ip` policy field now use the address your own proxy appended to `X-`

`Forwarded-For`, never an earlier caller-supplied entry. IP-range policies and per-IP rate limits (including device enrollment) now see one address per office or fleet rather than per workstation.

! Self-hosted deployments with more than one appending proxy must set `PRE_AUTH_TRUSTED_PROXY_HOPS` (Terraform: `trusted_proxy_hops`) to match; `CLIENT_IP_HEADER` is honoured only when it carries a single IP address.

! A tool call sending someone else's session ID can no longer block that session via 'Block later tool calls when any scanner blocks', add scan results or security events to that session, or mark an operator Stop as having taken effect in it. Hook events naming another session are not recorded into it.

! Runlayer now drops `X-Runlayer-User-Email`, `X-Runlayer-Identity-Token`, and other identity-forward headers sent by a client — any upstream server or workflow that relied on forwarding these caller-supplied headers will no longer receive them.

! `manual_oauth_client_secret` and `broker_vendor_client_secret` on `runlayer_server` are now marked sensitive in Terraform; module outputs that reference them must add `sensitive = true` or the apply will fail.

! MCP servers on Kubernetes that combine `enable_db` with an IRSA annotation (`eks.amazonaws.com/role-arn`) on the MCP ServiceAccount are now rejected by the operator (`ExtendDBReady=False` / `IRSAConflict`); remove one or the other to restore functionality.

WAS THIS USEFUL?  

EVALUATE

◆ AI Observability & Evals

ongrid

Feedback ⓘ v0.19.0 NOTES ↗ RANK

An ops AI Agent that understands your infrastructure, finds the root cause, and fixes it — right from Slack, Telegram, Lark or DingTalk.

> ongrid's v0.19.0 release adds scoped external MCP access with durable operations, surfaces stream client runtime state through the Channels API and UI, and introduces pagination for large switch interface queries in the network view.

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

01 **Stream client runtime state in Channels API and UI** NEW

34

The Channels API and UI now surface stream client runtime state, giving visibility into the live status of stream clients.

— Names API/UI surface but no endpoint, field, or exact path. v0.19.0

02 **Scoped external MCP access with durable operations** NEW

33

ongrid now supports scoped external MCP access paired with durable operations, letting the agent interact with external MCP servers under constrained permissions while operations persist reliably.

— No mechanism, limits, or config surface named. v0.19.0

03 **Paginated switch interface queries in network view** IMPROVED

32

ongrid now paginates large switch interface queries in the network view, improving handling of queries that return many interfaces.

— Names the area but no page size, flag, or endpoint. v0.19.0

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

h5i

Feedback ⓘ v0.5.2 NOTES ↗ RANK

Build, hack, and verify web apps. Find vulnerabilities through red-teaming.

> h5i v0.5.2 ships four plugin-based capabilities: Chromium browser session capture with request auditing, a websec workbench for replaying requests, a recon plugin for endpoint discovery, and CI-gated security regression testing against OpenAPI coverage.

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

01 **Browser session capture and request audit** NEW 90

`h5i browser open`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` capture, policy-check, and log every HTTP request made by a Chromium session. An allowlist is checked on every request and redirect hop, and `--in` adds an out-of-browser witness to the log.

Capture every request a Chromium session makes to a target, then review what was sent and what was blocked.

```
$ h5i browser open https://staging.example.com
# interact with the app, then:
h5i browser requests
h5i browser audit
```

— Four commands, allowlist mechanism, and runnable example all given. v0.5.2

02 **Recon plugin for endpoint discovery** NEW 80

Install the recon plugin with `h5i plugin install recon --from ./h5i-recon` to record endpoints, parameters, and candidate URLs a target exposes, grounded in stored messages.

Add the recon plugin after downloading its archive to discover endpoints exposed by a target.

```
$ curl -L https://github.com/h5i-dev/h5i/releases/download/v0.5.2/h5i-recon-v0.5.2-x86_64-unknown-linux-musl.tar.gz | tar -xz
h5i plugin install recon --from ./h5i-recon
```

— Install command plus a concrete download-and-run example. v0.5.2

03 **Security regression testing with coverage gating** NEW

80

Run repository-owned attack flows against a running server with `h5i test .h5i-tests/tests --target http://localhost:3000 --min-coverage 80` to fail CI if OpenAPI coverage drops below the given percentage.

Run repository-owned attack flows against a local server and fail CI if OpenAPI coverage drops below 80%.

```
$ h5i test .h5i-tests/tests --target http://localhost:3000 --min-coverage 80
```

— Exact CI command with flags and threshold value. v0.5.2

04 **Websec workbench plugin for request replay** NEW

70

Install the websec workbench with `h5i plugin install websec --from ./h5i-websec` to replay and manipulate HTTP requests through the same policy and budget as the browser.

— Install command given but no replay usage example shown. v0.5.2

WAS THIS USEFUL? 👍 👎

NoLabs nono

Feedback ⓘ v0.81.0 NOTES ↗ CODE ↗ RANK

The nono runtime runs AI agents in isolated environments with policy-controlled filesystem, network, and credential access.

> nono's v0.81.0 release adds a new policy-diagnostic command for Unix sockets, opt-in Linux namespace isolation, a new runtime-state policy group, Landlock V9 detection, structured proxy denial logging, and a breaking change requiring fully qualified profile names.

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

01 Unix socket policy diagnostics via `nono why` NEW

77

Query pathname Unix socket connect access with `nono why --unix-socket <path>` to diagnose `AF_UNIX` policy decisions, e.g. checking whether a sandboxed agent can connect to `/run/docker.sock`.

Check whether a sandboxed agent or tool is permitted to connect to a Unix socket such as the Docker daemon socket.

```
$ nono why --unix-socket /run/docker.sock
```

— Runnable command with example, but mechanism of decision logic not detailed v0.81.0

THINNER COVERAGE BELOW

02 `denial_category` field in proxy deny logs IMPROVED

55

Proxy deny log lines now include `denial_category`, giving operators a structured field to filter and alert on blocked network requests.

— Names the exact log field and its purpose for filtering/alerting v0.81.0

03 Fully qualified package profile names required BREAKING

50

Profiles now require fully qualified package profile names, making profile references unambiguous across registries; bare/unqualified profile name references that worked in v0.80.0 will break on upgrade.

— States the breaking behavior and upgrade impact but no migration command v0.81.0

04 Landlock ABI V9 detection IMPROVED

47

nono now detects Landlock ABI up to V9, extending future-proofing for Linux kernel upgrades while keeping enforcement at V6.

– Clear before/after on detection vs enforcement, but purely informational for users

v0.81.0

05 **Lower-overhead exec gate on Linux** IMPROVED 45

The exec gate on Linux now reads only the ELF regions it actually parses, reducing per-execution overhead and enabling lower-latency sandboxing for workloads that spawn many processes.

– Explains mechanism and benefit but is an automatic, non-actionable change v0.81.0

06 **`rtk_state` policy group** NEW 42

Add the `rtk_state` policy group to grant runtime-state access in profiles.

– Names the config key but gives no detail on what runtime-state access entails v0.81.0

07 **Opt-in namespace isolation (NEP-0004)** NEW 40

Enable opt-in namespace isolation on Linux (NEP-0004) for stronger per-tool containment.

– Names the feature and spec reference but no enabling flag or config shown v0.81.0

└──▷ **BREAKING ON UPGRADE**

! Profiles now require fully qualified package profile names; bare/unqualified profile name references that worked in v0.80.0 will break on upgrade.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 30, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue