

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 22, 2026 · EVERY DAY

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

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

▸ // HOW THIS ISSUE IS MADE

SHOW

VIEW

ISSUE VIEW

scroll



SCROLL

EXPAND
IN
PLACE

PREVIEW
PANE

INFORMATION DENSITY

everything



BRIEFED

EVERYTHING

Every tool in the issue, in full.

all 250 changes

Do you prefer **scroll**?



▸ // CONTENTS SHOW

70 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.

PROTECT

Make a leaked developer token worthless by morning

vim ? help

GitHub CLI's opt-in expiring OAuth credentials refresh automatically and work transparently with `git`, so the long-lived PAT sitting in a developer's config — the thing that turns one scraped laptop into months of repo access — stops being the default. 1Password adds post-quantum passkey support on Android in the same window.

[GitHub CLI](#) · [1Password](#)

PROTECT

Keep the endpoint agent installed when the attacker has admin

JumpCloud gates agent removal behind rotating maintenance tokens, closing the standard playbook step where an intruder with local admin uninstalls the agent and goes dark before anything is logged. Iru's FileVault key auto-rotation and phased OS update enforcement cut the manual chase of stale recovery keys and lagging patch cohorts.

[JumpCloud](#) · [Iru](#)

PROTECT

Hold a long analysis session instead of re-running it

OpenSandbox auto-pauses idle fast sandboxes and wakes them on access, with an SDK wake-retry contract so callers don't fail on the cold path — a detonation or agent run that would otherwise be torn down and restarted survives. Docker Sandboxes adds controls to block in-sandbox browser launches and OAuth redirects, closing a path where sandboxed code reaches out for real credentials.

[OpenSandbox](#) · [Docker Sandboxes](#)

OFFENSIVE SECURITY

ESP32 Marauder

Feedback



v1.19.1

NOTES ↗

RANK

A suite of WiFi/Bluetooth offensive and defensive tools for the ESP32.

ESP32 Marauder v1.19.1 extends hardware support to two new boards and reduces memory usage during wardriving on devices lacking PSRAM.

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

01 New hardware ports: M5Stack Core2 for AWS, Waveshare ESP32-C5-Touch-LCD-3.5

NEW

35

Added support for the M5Stack Core2 for AWS and Waveshare ESP32-C5-Touch-LCD-3.5 devices, letting Marauder run on these boards.

– Names two boards but no setup detail v1.19.1

02 Reduced no-PSRAM memory usage during wardriving IMPROVED

30

Reduces memory usage during wardriving on devices without PSRAM, improving reliability on lower-memory hardware.

– States the change but not the mechanism or amount v1.19.1

WAS THIS USEFUL?



◆ Reconnaissance & OSINT

nerva

[Feedback](#)[v1.77.0](#)[NOTES ↗](#)[RANK](#)

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

Nerva v1.77.0 extends its misconfiguration scanning to flag weak TLS cipher suite exposure.

└─> **WHAT SHIPPED** · **1 FEATURE** most completely described first [? what's the number?](#)

01 **Detection of weak TLS cipher suites** **NEW**

70

The `--misconfigs` scan now detects anonymous, NULL, and EXPORT TLS cipher suites as security misconfigurations, flagging weak TLS exposure on scanned targets.

Identify weak TLS cipher suite exposure on a target — anonymous, NULL, and EXPORT suites are now flagged as findings.

```
$ nerva -t example.com:443 --misconfigs --json
```

— Names exact flag and cipher types, but limited mechanism detail [v1.77.0](#)

[NEW VULNERABILITY CHECK](#)[MSSQL sa default blank password detection](#)[v1.76.0](#)[NOTES ↗](#)[STABLE](#)

└─> **GET THIS VERSION**

```
$ git clone --branch v1.76.0 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.76.0
```

WAS THIS USEFUL?



Hacktron

[Feedback](#)

Findings get an identity, a history, and a command palette

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

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

Hacktron's October release gives findings stable IDs with a full activity history and command-palette actions, adds a whitebox retest wizard to verify remediation on completed scans, and rounds out the findings list and whitebox scan experience with new filters, CWE tooltips, report refinements, and integration reliability fixes.

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

01 Finding identity, history, and command palette actions NEW

85

Findings now get human-readable, stable IDs shown in the row, panel, full page, and URL with click-to-copy. The finding detail panel adds an **Activity** tab showing the full detection and edit history with actor and changed fields per entry, switchable to a **Comments** view for just the thread. When a finding is open, the command palette exposes finding-scoped actions: change status, change severity, mark/unmark duplicate, create ticket, start remediation, copy ID, copy link, and open full page. Multi-ID search lets you paste two or more finding IDs into the search box to filter to exactly those findings.

Filter the findings list to a specific set of findings by pasting their IDs — useful when a ticket references multiple findings you want to triage together.

📌 In the findings search box, paste: HT-1042 HT-1078 HT-1091

— Names every UI surface and action, with a runnable search example.

Findings get an identity, a history, an...

02 Whitebox retest wizard for completed scans NEW

75

A new retest wizard lets you request a retest for a completed scan directly from the scans list or scan detail page, covering all reported findings except false positives and accepted risk. Retest status and results appear on the scan detail page once the retest completes.

Request a whitebox retest after a scan is complete to verify that reported findings (excluding false positives and accepted risk) have been remediated.

📌 Navigate to Scans > [completed scan] > Retest, then follow the retest wizard to submit the request. Track status and results on the scan detail page.

– Clear scope and exact navigation steps to trigger it.

Findings get an identity, a history, an...

03 Findings list filters and CWE tooltip NEW 60

Adds **Source**, **Reopened**, and **Date range** filters to the findings list, each saved in the URL, plus a CWE name tooltip on the CWE badge so the full weakness name is visible without leaving the page.

– Names each filter but no deeper mechanism given.

Findings get an identity, a history, an...

THINNER COVERAGE BELOW

04 GitHub.com integration failure reasons for on-prem appliances 55

IMPROVED

On-prem appliances that cannot reach **GitHub.com** now disable the **GitHub.com** integration card with a specific reason (DNS, blocked, or TLS) instead of failing at the code exchange step.

– Names failure modes but no reader action described.

Findings get an identity, a history, an...

05 Whitebox scan report methodology and coverage detail updates 50

IMPROVED

The **Methodology** section of a whitebox scan report now shows how many functions each area selected instead of an arbitrary list of five names, and **Coverage Detail** caps the unselected entry-points table at 50 rows.

– Specific before/after but limited actionable follow-up.

Findings get an identity, a history, an...

06 Inline browser-auth audit instructions in scan wizard 50

IMPROVED

The target step of the whitebox scan wizard now shows browser-auth audit instructions directly on the page with live verification status, replacing the previous separate modal.

– Describes the change but not the verification mechanism.

Findings get an identity, a history, an...

07 Trial/subscription notice on Repositories page NEW 45

The **Repositories** page shows a notice when PR Review is toggled on for a repo but no active trial or subscription backs it, with a direct link to start a trial.

– Names the UI location and link but little mechanism.

Findings get an identity, a history, an...

WAS THIS USEFUL?



APPLICATION SECURITY

♦ WAF, Bot & API Protection

Akto

Feedback



announced · Oct 22, 2026

BLOG ›

Akto discovers APIs, MCP servers and AI agents, tests and red-teams them, and enforces runtime guardrails and posture management.

Akto is now a native LiteLLM guardrail provider with `pre_call` / `post_call` modes, `default_on`, and a reference architecture at github.com/akto-api-security/akto-litellm.

NEW INTEGRATION

NEW DETECTION

Adds `guardrail: akto` as a native LiteLLM guardrail provider, configured via `guardrail_name`, `akto_base_url` (`os.environ/AKTO_GUARDRAIL_API_BASE`), and `akto_api_key` (`os.environ/AKTO_API_KEY`) in `config.yaml`.

Supports `mode: pre_call` and `mode: post_call` guardrail entries for Akto, covering PII scanning on both inbound user prompts and outbound model responses independently.

Enables `default_on: true` on individual Akto guardrail entries in `config.yaml` to enforce PII scanning across all traffic without per-request opt-in.

Provides a reference architecture for a Claude Code → LiteLLM → Bedrock → Anthropic deployment with layered Akto guardrails at github.com/akto-api-security/akto-litellm.

WAS THIS USEFUL?



Open-source and cloud-native Web Application Firewall (WAF)

BunkerWeb 1.6.17 adds configurable database connect/read timeouts, bad-behavior report size caps, a shared-memory zone size knob for dynamic upstreams, Let's Encrypt certificates for public IP addresses with short-lived profiles, automatic ACME HTTP-01 challenge passthrough, and an expanded default Permissions-Policy header.

↳ **WHAT SHIPPED** • **6 FEATURES** most completely described first [? what's the number?](#)

01 Configurable database connect and read timeouts **NEW**

95

Adds `DATABASE_CONNECT_TIMEOUT` to set the maximum wait when opening a connection to a MySQL, MariaDB, or PostgreSQL server, and `DATABASE_READ_TIMEOUT` to set the maximum wait for a response from a MySQL or MariaDB server before failing (PostgreSQL relies on TCP keepalives instead). Both accept a time suffix (ms, s, m, h, d, w, M, y) and are set in `/etc/bunkerweb/variables.env`.

Prevent a stale MySQL connection from hanging the scheduler indefinitely by capping connect and read waits.

```
env

DATABASE_CONNECT_TIMEOUT=5s
DATABASE_READ_TIMEOUT=30s
```

Set database connection and read timeouts so a lost MySQL peer does not hang BunkerWeb indefinitely.

```
ini

DATABASE_CONNECT_TIMEOUT=5s
DATABASE_READ_TIMEOUT=10s
```

Set a 5-second connection timeout and a 30-second read timeout for a MariaDB backend so a vanished peer never blocks BunkerWeb indefinitely.

```
ini

DATABASE_CONNECT_TIMEOUT=5s
DATABASE_READ_TIMEOUT=30s
```

Reduce the database connect timeout in a high-availability environment where a dead DB node should be detected and bypassed quickly.

```
ini

DATABASE_CONNECT_TIMEOUT=3s
DATABASE_READ_TIMEOUT=1m
```

– Named config keys, file path, time units and runnable examples [v1.6.17](#)

02 Let's Encrypt certificates for IP addresses and profiles NEW

95

Supports Let's Encrypt certificates for public IPv4 addresses: add the IP alongside hostnames in the server name field (e.g. `SERVER_NAME="www.example.com 203.0.113.10"` with `MULTISITE=no`) and BunkerWeb requests a single certificate covering all entries. IP certificates use the `shortlived` certificate profile, have 160-hour validity, support up to 25 names, and are auto-renewed daily.

Request a Let's Encrypt certificate that covers both a hostname and a public IPv4 address in single-site mode.

```
ini

SERVER_NAME="www.example.com 203.0.113.10"
MULTISITE=no
```

– Names the config fields, profile name and exact limits [product docs](#)

03 Bad-behavior report size limits NEW

90

Adds `BAD_BEHAVIOR_REPORT_MAX_EVENTS` to cap the number of new events copied as ban evidence into bad-behavior reports, and `BAD_BEHAVIOR_REPORT_MAX_STRING_SIZE` to cap the maximum byte length of each URL or method string stored in those reports, controlling report storage growth under high-volume attack traffic.

Limit how much memory bad-behavior ban reports consume when logging high-volume attack traffic.

```
env

BAD_BEHAVIOR_REPORT_MAX_EVENTS=50
BAD_BEHAVIOR_REPORT_MAX_STRING_SIZE=512
```

Cap ban-evidence report size to control storage growth when bad-behavior events are frequent.

```
ini
```

```
BAD_BEHAVIOR_REPORT_MAX_EVENTS=50  
BAD_BEHAVIOR_REPORT_MAX_STRING_SIZE=512
```

Cap ban-report evidence to the 20 most recent events, each URL or method stored up to 512 bytes, to control report storage growth.

```
ini
```

```
BAD_BEHAVIOR_REPORT_MAX_EVENTS=20  
BAD_BEHAVIOR_REPORT_MAX_STRING_SIZE=512
```

– Both config keys named with concrete example values [product docs](#)

04 ACME HTTP-01 challenge passthrough IMPROVED 70

ACME HTTP-01 challenge URIs now pass through automatically, bypassing antibot, blacklist, rate limiting, Basic Auth, ModSecurity, and the ban check, so Let's Encrypt validation succeeds even with those protections active.

– Mechanism and scope named but no config toggle shown [product docs](#)

THINNER COVERAGE BELOW

05 Shared memory zone size for dynamic upstreams NEW 50

Adds `LOADBALANCER_ZONE_SIZE` config key to control the shared memory zone size for upstreams with dynamic resolution.

– Key named but no default value or example given [product docs](#)

06 Expanded default Permissions-Policy header IMPROVED 45

Expands the default `PERMISSIONS_POLICY` header to also deny `unbounded-element`.

– Names the header and directive but gives no further context [v1.6.17](#)

↳ ALSO FROM THESE RELEASES

•
•
•

WAS THIS USEFUL?



DataDome protects web, mobile, and API applications by detecting and blocking bot and fraud traffic.

DataDome's biggest addition this window is a full server-side and client-side bot-protection integration for the Azion edge platform, covering firewall-function arguments, GraphQL/MCP protocol support, and JS Tag injection, alongside a smaller Fraud SDK Java update.

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

01 GraphQL and MCP protocol support in Azion firewall

NEW

90

The firewall function adds `enableGraphQLSupport` and `graphqlEndpoint` to extract GraphQL operation name and type on GraphQL requests, and `enableMCPSupport` and `mcpEndpoint` to enable Model Context Protocol support and extract MCP information at a configurable endpoint. Both extractions share a `maximumBodySize` cap on the request body read, defaulting to `25600` bytes (25 KB).

Enable Model Context Protocol support on a custom MCP endpoint and raise the body-read limit for large MCP payloads.

```
json
{
  "serverSideKey": "<your-server-side-key>",
  "enableMCPSupport": true,
  "mcpEndpoint": "/mcp,/mcp-v2",
  "maximumBodySize": 51200
}
```

– Mechanism, defaults and args named with runnable example

[product docs](#)

02 Azion firewall function arguments for bot-protection control

NEW

85

The `firewall.js` Azion firewall function gains server-side arguments `serverSideKey`, `endpoint`, `timeout`, `urlPatternExclusion`, `urlPatternInclusion`, `staticAssetExclusion`, `encodedCharInclusion`, `removeMatrixParams`, `ipExclusion`, `logEnrichedHeaders`, `enableDebugging`, `customFieldString1-3`, `customFieldInteger1-2`, `customFieldFloat1`, `userID`, and `productID` for fine-grained control over which requests are inspected and how.

Scope bot-protection to login and checkout paths only, exclude your admin area, and enable GraphQL extraction for your `/graphql` endpoint — all in the firewall function Arguments.

```
{
  "serverSideKey": "<your-server-side-key>",
  "urlPatternInclusion": "/login|/checkout",
  "urlPatternExclusion": "/admin/",
  "enableGraphQLSupport": true,
  "graphqlEndpoint": "/graphql"
}
```

– Names every argument with a runnable config example [product docs](#)

03 JS Tag injection arguments for Azion NEW INTEGRATION

NEW

85

The `js-tag-injection.js` Azion application function adds client-side arguments `clientSideKey`, `jsURL`, `jsTagOptions`, `originURL`, and `maximumResponseBodySize` to inject the DataDome JS Tag into HTML responses. `maximumResponseBodySize` (default 2 MB) sets the HTML response size threshold above which the JS Tag is skipped and the response passed through unmodified.

Inject the DataDome JS Tag via your own first-party domain and raise the response-body buffer limit to 4 MB for large HTML pages.

```
{
  "clientSideKey": "<your-client-side-key>",
  "jsURL": "https://ddfoo.example.com/tags.js",
  "originURL": "https://origin.example.com",
  "maximumResponseBodySize": 4194304
}
```

– Named args, default threshold, and runnable config example [product docs](#)

THINNER COVERAGE BELOW

04 Effective event time field in Fraud SDK Java NEW

25

Fraud SDK Java 3.1.0 adds an optional `eventDate` field to specify the effective event time when submitting fraud events.

– Only the field name is given, no example or mechanism [Fraud SDK Java 3.1.0 \(2026-10-21\)](#)

NEW INTEGRATION [Fastly Compute \(Rust\) integration](#)

Fastly Compute Rust module 1.0.0 (2026-10-21)

[BLOG ↗](#)

LATEST

STABLE

Azion module 1.0.0 (2026-10-21)

[BLOG ↗](#)

STABLE

WAS THIS USEFUL?



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

Fastly's Terraform Provider 10.0.0 adds multi-service monitoring to the alert resource, but breaks configurations that set `user_id` on the AI runtime control virtual key resource.

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

01 `'user_id'` now computed in AI runtime control virtual key

70

BREAKING

The `user_id` argument in `ai_runtime_control_virtual_key` is no longer configurable — it is now a computed (read-only) attribute. Any Terraform configuration that sets `user_id` will fail to plan after upgrading.

— Names exact resource and field plus plan failure impact

Fastly Terraform Provider 10.0.0 (2026-...

THINNER COVERAGE BELOW

02 Multi-service monitoring in alert resource

NEW

45

The Terraform provider's `alert` resource now supports monitoring multiple services from a single resource definition, rather than requiring one alert resource per service.

— Names the resource but not config syntax or limits

Fastly Terraform Provider 10.0.0 (2026-...

↳ BREAKING ON UPGRADE

! The `user_id` argument in `ai_runtime_control_virtual_key` is no longer configurable — it is now a computed (read-only) attribute. Any Terraform configuration that sets `user_id` will fail to plan after upgrading.

WAS THIS USEFUL?



Semgrep

Feedback



changelog · Oct 22, 2026

CHANGELOG >

RANK

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

Semgrep extended Agentic Workflows run limits and improved credit-usage reporting on the workflow scheduling page.

L--> WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Credits used column counts all workflow triggers IMPROVED

40

The **Total credits used** column on the workflow scheduling page now counts runs from every trigger, including pull-request-initiated runs, not only scheduled runs.

– Names exact UI column and scope of change. [changelog-20261022-d485ed84](#)

02 Agentic Workflows run time extended to five hours IMPROVED

30

Agentic Workflows runs can now take up to five hours instead of failing at the previous two-hour limit.

– States the before/after limit but no mechanism. [changelog-20261022-d485ed84](#)

L--> ALSO FROM THESE RELEASES

Retrieve Autofix PR and failure details for every finding in a deployment – useful for auditing which findings have pending fix PRs or repeated Autofix failures.

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

Rotate a stale GitHub App webhook secret so that pull request scans resume triggering after secret rotation.

📌 In Semgrep AppSec Platform, go to Settings > Source code managers, find your GitHub App connection, and click **Rotate webhook secret**.

WAS THIS USEFUL?



HashiCorp terraform

[Feedback](#)[v1.17.0](#)[NOTES ↗](#)[RANK](#)

Terraform enables you to safely and predictably create, change, and improve infrastructure. It is a source-available tool that codifies APIs into declarative configuration files that can be shared amongst team members, treated as code, edited, reviewed, and versioned.

Terraform v1.17.0 promotes policy enforcement in `plan`, `apply`, and `query` to GA, adds a `-minimal-refresh` flag to speed up plan/apply runs, and allows variables and locals in provider requirements blocks.

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

01 GA policy enforcement across plan, apply, query NEW

86

The `-policies` flag on `plan`, `apply`, and `query` is now GA and no longer requires an experimental build or `-allow-experimental-features`. Query policies evaluate resources found by `list` blocks and return human-readable or JSON results.

Enforce org-wide policies in a standard (non-experimental) pipeline by running plan with the now-GA `-policies` flag.

```
$ terraform apply -policies ./policies/ -auto-approve
```

– Names flag, affected commands, and JSON/human output modes [v1.17.0](#)

02 `-minimal-refresh` planning flag NEW

80

Adds `-minimal-refresh` to `plan` and `apply`, refreshing only resources that have proposed changes rather than reading the full state, which can speed up CI plan runs by skipping unchanged remote state reads.

Speed up CI plan runs by only refreshing the resources that are actually changing, skipping unchanged remote state reads.

```
$ terraform plan -minimal-refresh -out=tfplan
```

– Includes exact flag and a runnable command [v1.17.0](#)

THINNER COVERAGE BELOW

03 Variables and locals in provider requirements NEW

50

Supports variables and locals in provider `required_providers` blocks, enabling dynamic provider version and source constraints.

– Clear mechanism but no concrete example given `v1.17.0`

04 **format_version field in terraform version JSON** IMPROVED

 50

Adds a `format_version` field to the JSON output of `terraform version`, enabling safer future evolution of that output format.

– Names field and purpose but no example `v1.17.0`

05 **mock_provider support for ephemeral resources in tests** NEW

 45

Adds `mock_provider` support for ephemeral resources in `terraform test`.

– Named surface but no usage detail `v1.17.0`

WAS THIS USEFUL?



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

GitHub CLI added opt-in, expiring OAuth tokens that automatically rotate and become useless within hours if leaked, with full support for setting them as default, checking expiration, and using them transparently with git.

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

01 Expiring, auto-refreshing OAuth tokens in GitHub CLI NEW

93

GitHub CLI v2.103.0 adds `--expiring` flag to `gh auth login` to request short-lived OAuth tokens that automatically rotate and become useless within hours if leaked. Running `gh config set expiring_credentials enabled` makes expiring tokens the default for every future sign-in without passing `--expiring` each time, and `gh auth status` now shows the expiration time of the current access token when expiring credentials are active. `gh auth setup-git` configures gh as the git credential helper so HTTPS git operations get automatically refreshed tokens. Support requires GitHub Enterprise Server 3.23+, falling back to non-expiring behavior on earlier versions; scripts that cache the output of `gh auth token` instead of reading it fresh each time will eventually hold an expired token.

Opt in to expiring tokens for a single login so a leaked access token expires within hours rather than lasting indefinitely.

```
$ gh auth login --expiring
```

Make expiring tokens the default for every future sign-in on this machine without passing `--expiring` each time.

```
$ gh config set expiring_credentials enabled
```

Configure gh as the git credential helper so expiring tokens are served and refreshed automatically for HTTPS git operations.

```
$ gh auth setup-git
```

– Names exact flags, config key and commands with mechanism and version scope

GitHub CLI: Opt in to expiring, auto-re...

➡ BREAKING ON UPGRADE

! Scripts or extensions that run `gh auth token`, store the result, and reuse it later will eventually hold an expired token — read it fresh each time with `gh auth token` instead of caching it.

WAS THIS USEFUL?



Snyk Open Source

Feedback



changelog · Oct 21, 2026

BLOG ↗

RANK

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.

Snyk Open Source's CLI now validates that `--reachability=true` is set when using `--source-dir`, catching a silent misconfiguration.

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

01 Validation for `--source-dir` with reachability flag IMPROVED

70

`snyk test --source-dir` now checks that `--reachability=true` is also set, returning a clear error instead of silently ignoring the flag or misinterpreting the path.

Run a reachability-aware scan against a specific source directory so Snyk validates the flag combination and surfaces reachable vulnerabilities.

```
$ snyk test --source-dir=./src --reachability=true
```

– Names exact flags and command; limited to a single validation fix 2026-10-21-6a37e478

WAS THIS USEFUL?



Cisco skill-scanner

[Feedback](#)[2.3.1](#)[NOTES ↗](#)[CODE ↗](#)[RANK](#)

The skill-scanner is a security scanner for AI agent skills

Skill Scanner 2.3.1 adds programmatic lookup of the full AITech taxonomy and refreshes stale LLM provider defaults, with two breaking changes for teams pinning taxonomy hashes or default models.

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

01 AITech taxonomy definition lookup and reference data

NEW

78

Adds `get_aitech_definition()` and `get_aisubtech_definition()` functions to look up the bundled definition for any AITech objective, technique, or sub-technique code directly from the snapshot. Ships a generated taxonomy reference covering all 19 objectives, 44 techniques, and 130 sub-techniques, with each entry in `threats.py` carrying a full description and AI Defense labels.

Look up the human-readable definition of an AITech code surfaced in a finding, without leaving Python.

```
python

from skill_scanner import get_aitech_definition,
get_aisubtech_definition

print(get_aitech_definition('AITech-1.1'))
print(get_aisubtech_definition('AITech-1.1.2'))
```

– Names exact functions, file, and taxonomy counts with a runnable example [2.3.1](#)

02 LLM provider default model updates

BREAKING

78

Updates the `--llm-provider openrouter` default model from `openrouter/openai/gpt-4` to `openrouter/openai/gpt-5.6-sol`, and the `gcp-vertex` default from `vertex_ai/gemini-1.5-pro` to `vertex_ai/gemini-3.8-flash`, replacing retired models. Users relying on either default now get the new model automatically; set `SKILL_SCANNER_LLM_MODEL` to pin a specific model instead.

– Gives exact old/new model names and the env var to pin [2.3.1](#)

THINNER COVERAGE BELOW

03 Taxonomy snapshot hash changed

BREAKING

45

`scan_metadata.taxonomy.snapshot_sha256` changes from `3df746f0...` to `413e8ce9...` ; anything that pins this hash must be updated.

– Names exact field and hash values but no further context 2.3.1

↳ **BREAKING ON UPGRADE**

! `scan_metadata.taxonomy.snapshot_sha256` changes from `3df746f0...` to `413e8ce9...` ; anything that pins this hash must be updated.

! OpenRouter users relying on the default model now get `openrouter/openai/gpt-5.6-sol` instead of `openrouter/openai/gpt-4` ; set `SKILL_SCANNER_LLM_MODEL` to pin a specific model.

! Vertex AI users relying on the default model now get `vertex_ai/gemini-3.8-flash` instead of `vertex_ai/gemini-1.5-pro` ; set `SKILL_SCANNER_LLM_MODEL` to pin a specific model.

WAS THIS USEFUL?



cve-lite-cli

[Feedback](#)[v1.40.0](#)[NOTES ↗](#)[RANK](#)

Fast, developer-friendly JS/TS dependency vulnerability scanner with local lockfile scanning, OSV matching, direct vs transitive visibility, --fix, JSON output, and practical remediation guidance.

cve-lite-cli v1.40.0 adds PR-scoped delta findings comments to its GitHub Action and enriches JSON output with run metadata, alongside three breaking fixes to ratchet gating, pnpm v6 lockfile parsing, and transitive fix-version resolution.

←→ **WHAT SHIPPED · 6 FEATURES** most completely described first [? what's the number?](#)

01 --ratchet gating without committed baseline now fails

90

BREAKING

--ratchet with no committed baseline now applies --fail-on and exits non-zero when findings match the gate, instead of writing a baseline and exiting 0 — a CI job that checks out fresh every run and never committed its baseline could previously not fail and now can.

– Full before/after mechanism with exact flags named [v1.40.0](#)

02 PR-scoped delta comments in GitHub Action **NEW**

80

Adds a delta input to the GitHub Action so pull requests receive a single upserted comment listing only the findings that PR introduced, not the full dependency tree; fork PRs fall back to the job summary instead.

– Names exact input and behavior, missing config example [v1.40.0](#)

03 pnpm lockfile v6 parsing fix

75

BREAKING

pnpm lockfile v6 (pnpm 8 and earlier) was misparsed, causing vulnerable dependency trees to report clean; re-scanning after upgrading may surface new findings that were silently missed before.

– Names affected pnpm versions and impact, no repro steps [v1.40.0](#)

04 Transitive finding upgrade recommendation resolves correct parent

60

BREAKING

The recommended upgrade version for a transitive finding now resolves the parent from the finding's own dependency path, so some previously reported fix versions will change.

– States behavior change but no concrete example given v1.40.0

THINNER COVERAGE BELOW

05 Enriched `--json` output metadata IMPROVED 50

`--json` output now includes version, timestamp, and effective options alongside findings.

– Names flag and fields but no schema or example shown v1.40.0

06 SARIF rule help links to project IMPROVED 40

SARIF rule help markdown now leads with a link to the project so a triaging analyst can identify which tool produced each Code Scanning alert.

– Describes change but lacks concrete format or path detail v1.40.0

! --> BREAKING ON UPGRADE

- ! `--ratchet` with no committed baseline now applies `--fail-on` and exits non-zero when findings match the gate, instead of writing a baseline and exiting 0 – a CI job that checks out fresh every run and never committed its baseline could previously not fail and now can.
- ! pnpm lockfile v6 (pnpm 8 and earlier) was misparsed, causing vulnerable trees to report clean; re-scanning after upgrading may surface new findings that were silently missed before.
- ! The recommended upgrade version for a transitive finding now resolves the parent from the finding's own dependency path, so some previously reported fix versions will change.

WAS THIS USEFUL?



Phoenix Security

[Feedback](#)[Integration with Microsoft Azure Connector \(2026-10-21\)](#)[BLOG](#)[RANK](#)

Phoenix Security correlates scanner, cloud and runtime findings with exploitability and ownership, and uses AI agents to prioritise and open remediation pull requests.

Phoenix Security shipped a new Microsoft Azure connector that lets teams pull AKS and Container Apps runtime inventory, Container Registry image data, and Entra ID group membership into its correlation and remediation workflow.

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

01 Microsoft Azure connector with AKS, Container Apps, ACR and Entra ID sync

NEW

95

A new Microsoft Azure connector under [Integrations > Connectors](#) authenticates via a customer-managed Microsoft Entra service principal using Tenant ID, Client ID, and Client Secret. Four capabilities can be independently enabled, validated, renamed or removed without deleting the base connection: Azure Kubernetes Service runtime (imports running-container context for AKS plus Defender for Cloud security resource data, requiring [Reader](#) and [Security Reader](#) roles on each scoped subscription), Azure Container Apps runtime inventory ([Reader](#) role on each scoped subscription), Azure Container Registry image inventory (imports repository, image digest and manifest inventory using [Reader](#) plus data-plane read access), and optional team synchronisation from Microsoft Entra ID groups using [Microsoft Graph](#) application permissions [GroupMember.Read.All](#) and [User.Read.All](#) with tenant admin consent.

Set up the base Azure connection and enable AKS runtime discovery so Phoenix Security can pull running-container context from your Kubernetes clusters.

📌 In Phoenix Security, go to Integrations > Connectors and select Microsoft Azure. Enter a connector name, your Tenant ID, Client ID, and Client Secret, then select Create. Once created, open the connector from the Connected connectors list, select the Kubernetes Service capability, keep Enabled on, and select Save.

Enable Entra ID team synchronisation so Phoenix team membership stays current from Microsoft Entra ID groups, using least-privilege Graph permissions.

📌 In the Microsoft Entra admin centre, add application permissions [GroupMember.Read.All](#) and [User.Read.All](#) to the Phoenix Security Azure Connector app registration and grant tenant admin consent. Then in Phoenix

Security under Integrations › Connectors, open the Azure connector, select the Team sync from Entra ID groups capability, keep Enabled on, and select Save.

– Names exact UI path, auth fields, permissions and roles per capability.

Integration with Microsoft Azure Connec...

WAS THIS USEFUL?



Black Duck Feedback ⓘ

Polaris release update: Deeper detection, remediation you control, scanning wherever you

BLOG ↗ RANK

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

Black Duck's Polaris update brings AI-framework-aware static analysis, finer control over automated fix PRs, container scanning triggerable from more CI/CD surfaces, and support for GitHub Enterprise Cloud data residency.

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

01 Per-assessment-type fix PR control IMPROVED 83

Adds per-assessment-type fix PR control in the CLI and Polaris UI, letting teams enable automated remediation for SCA only, Static only, or both independently. Extends Static AI fix PRs to the Jenkins plug-in and Black Duck Security Setup workflows, which previously supported SCA fix PRs only; enabling fix PRs in these two surfaces still turns on both SCA and Static issues together.

– Explains mechanism and before/after across named surfaces

Polaris release update: Deeper detectio...

02 Container scan automation and layer-level risk analysis NEW 82

Adds container scan initiation from GitHub Actions, GitLab templates, Jenkins plug-in, Bitbucket pipe, ADO extension, the CLI, and Black Duck Security Setup workflows, with layer-level risk analysis identifying which container layer introduced a specific vulnerability or component. A unified component inventory combines container scan findings with binary, package manager, and signature scan results into one view.

– Many named entry points and a clear mechanism, no exact syntax

Polaris release update: Deeper detectio...

03 AI framework static analysis in Rapid Scan Static NEW 73

Rapid Scan Static (Sigma 2026.9.1) adds specialized static security analysis for AI frameworks — the Anthropic SDK, OpenAI SDK, and LangChain — catching dangerous tool enablement, agent configuration risks, and AI prompt injection issues.

– Names exact frameworks and engine version but no run command

Polaris release update: Deeper detectio...

04 **Rapid Scan Static language coverage and accuracy** IMPROVED

73

Extends Rapid Scan Static language coverage to Go 1.27, expands detection across Groovy (including Grails and Gradle), JavaScript, and Python, improves detection accuracy for CORS and TLS misconfigurations in Groovy/Grails applications, and reduces false positives in JavaScript unit test files.

– Specific languages and fixes named, no exact command given

Polaris release update: Deeper detectio...

05 **GitHub Enterprise Cloud data residency support** NEW

60

Adds support for GitHub Enterprise Cloud (GHEC) tenants with data residency, allowing connection to customer-specific regional GitHub hosts with automatic regional API routing via the CLI, GitHub Actions, or the Polaris UI.

– Names surfaces and mechanism but lacks setup steps

Polaris release update: Deeper detectio...

ALSO FROM THESE RELEASES

.
. .

WAS THIS USEFUL?



Aikido Security

Feedback



announced · Oct 20, 2026

[BLOG](#)

RANK

Aikido Security is a developer-focused application security platform that scans code, dependencies, containers, and cloud infrastructure to detect and help fix vulnerabilities.

Aikido launched Noon, a brand-new workforce security platform separate from its developer-focused scanning tools, spanning email security, device protection, AI governance, and phishing/vishing/deepfake training simulations.

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

01 Noon phishing, vishing and deepfake training simulations NEW

62

Noon's training and simulations module delivers realistic phishing, vishing, smishing, and deepfake attack scenarios personalized by role, behavior, and risk, with lessons delivered in Noon or directly in Slack, Teams, or Google Chat.

– Names attack types and delivery channels, giving a clear starting point

launch-20261020-1d93175b

THINNER COVERAGE BELOW

02 Noon device protection module NEW

45

Noon's device protection module covers employee devices against malicious packages, installs, extensions, agent skills, AI tools, risky websites, and apps.

– Lists concrete threat categories but no setup or config detail

launch-20261020-1d93175b

03 Noon workforce security platform launch NEW

38

Aikido introduced Noon, a new workforce security platform distinct from Aikido's existing developer-focused tooling, combining email security, device protection, AI governance, and phishing training into one product.

– Announcement names four product areas but no mechanism for AI governance

launch-20261020-1d93175b

04 Noon email security module NEW

31

Noon's email security module autonomously detects and removes inbox threats and malicious emails that traditional filters miss.

– States the behavior but not how detection works

launch-20261020-1d93175b

NEW INTEGRATION

[Aikido AI Pentest integration with Drata](#)

announced · Oct 19, 2026

[BLOG](#) ↗

Aikido AI Pentest findings and fix evidence now sync automatically into Drata, with 6 months of unlimited retesting included.

Aikido AI Pentest findings and compliance evidence sync directly into Drata, eliminating manual spreadsheet exports for audit prep.

Every Aikido AI Pentest through the Drata offer includes 6 months of unlimited retesting, so fixes can be verified the same week they ship.

WAS THIS USEFUL?



OpenQodex

[Feedback](#)[openqodex@0.13.0](#)[NOTES](#)[RANK](#)

Open source code review that runs in your coding agent before you push. Scanners plus a separate AI reviewer, on your own login, no key.

OpenQodex built out a full code graph for change-impact analysis this window — now interface-aware, cross-package, persistently cached, and exposed via a new `openqodex graph` command and MCP server — alongside a `reviewChange` library API, branch/PR review targets, a hardened GitHub Action with Claude Code AI review, scanner installation verification, and a consolidated single-command review flow with completeness checks.

WHAT SHIPPED • 23 FEATURES most completely described first ? what's the number?

01 Library API for programmatic review (`reviewChange`)

NEW

93

Adds a `reviewChange(input, reviewer, options)` library function that reviews a change using the host's model in server profile, running five read tools over a head-commit snapshot and returning findings, coverage, a model completion record (contract `openqodex-model-review-1`), and a report in markdown, SARIF, and JSON. Options include `previousReviewedSha` (review only what changed since the last review while still checking findings on the whole change), `scopes` (keeps change/snapshot/scanners/graph/tools/context inside named folders), `context` (supplies lessons, comments, summaries, notes, and earlier findings as quoted data), `secondReviewer` (runs a second model on the same brief and merges both models' findings), `authorize` (a budget callback with `requestChars`, called before every model call, with no call starting after `deadlineMs`), and `instructions` (passes repository owners' rules quoted to the model), with each finding now carrying a `suggestedChange`. It was later extended with `result.rounds` (also saved as `rounds.json` beside the report, recording failed checks, the earlier answer, and changed line ranges per correction round), `result.scan` (the full redacted scan record including candidates, scanner summaries, rule packs, `selection`, and `trail`), and `result.timings` so server hosts can read timings without a file. The library also exports `isFixturePath` and `ruleClassFor`, an import surface for scanners, code graph, lenses, config parser, and renderers with types, and a lens marked `security: true` now carries that flag when loaded.

— Full function signature, options, and return fields named. [openqodex@0.13.0](#)

[openqodex@0.12.0](#)

02 Code graph for change-impact analysis

NEW

98

OpenQodex adds a code graph to the review brief for TypeScript, JavaScript, Python, Go and Ruby that shows which functions a change touches, exact call-site lines, file-import

relationships, and removed functions still referenced elsewhere, controlled by `--no-graph` and config keys `graph.enabled`, `graph.budget_ms`, `graph.max_files`, and `graph.max_file_bytes` (results cached per file under `.openqodex/graph/`). It was extended with cross-package caller resolution across pnpm/npm/yarn workspaces (using each file's nearest `tsconfig.json` and Python `src` roots), caller confidence levels `certain`, `likely`, and `possible` each with a named cause, a persistent cache governed by `graph.max_cache_mb` (default 512 MB) and `graph.max_heap_mb` (with `graph.max_files` now counting only new parses rather than all files), kept reachable via a local git ref `refs/openqodex/graph/<tree>`, graph execution on manifest-only changes (`package.json`, `tsconfig.json`, `pyproject.toml`, `go.mod`, ...), a snapshot under `.openqodex-review/graph/` with truncation limits reported, hidden subcommands `openqodex graph build | status | search | symbol | callers | callees | importers | changes | unknowns | explain | capabilities` (each with `--json`, documented in `docs/graph.md`), interface-aware call tracing through interfaces, abstract classes, and base classes using Python MRO, Go embedding/method sets, and Ruby prepend/include/extend order (interface calls capped at 32 possible implementations), possible-caller tracking for functions passed as parameters, locally assigned functions, literal-table entries, and functions returned by name (capped at 20 inline), and `implementers/` and `references/` pages per touched symbol added to the review packet with `unknowns.json` extended to list framework gaps.

– Mechanism, config keys, and limits fully named across releases. `openqodex@0.11.0`

`openqodex@0.10.0`

`openqodex@0.2.0`

03 Scanner installation and verification

NEW DETECTION

BREAKING

85

Scanner downloads are now scoped per project — brakeman, rubocop, oxlint, ruff, and shellcheck only activate when their manifests or markers are detected. osv-scanner was updated to 2.6.0, adding support for `bun.lock`, `uv.lock`, `pdm.lock`, `pylock.toml`, and NuGet lockfiles; oxlint was updated to 1.86.0 and now installs from its GitHub release verified against its sha256 with no npm dependency; semgrep, bandit, brakeman, and rubocop now install from lock files shipped in the package with every dependency pinned to a version and sha256; brakeman now requires Ruby 3.0 or newer. `doctor --install` is now scoped to a repository's actual file needs (previously installed everything) — `doctor --install --all-scanners` restores the old behaviour, and `doctor --install --all-scanners --require-all` verifies every scanner is installed at its pinned version with its runtime, printing one line per scanner that is not ready and exiting with code 2 if any is missing. The same check is available to library consumers via `preinstallScanners({ installRoot, require })` and `createToolResolver({ installRoot })`, and `runScanners` gained a `scratchRoot` parameter so concurrent server runs write only inside their own folder.

– Commands, versions, and flags all concretely named. `openqodex@0.12.0` `openqodex@0.9.0`

Adds `openqodex graph` as a menu command answering questions about callers, callees, importers, implementers, overrides, value/type uses, connection paths, reachability walks, file/folder definitions, project dependencies, and import cycles (`openqodex graph help` lists all questions), with `--budget-ms` (default 1 second), `--limit` (default 50 items), `--cursor`, and `--tokens` flags. Adds `openqodex mcp` to serve the same code-graph questions to coding agents as MCP tools over stdio, including `graph_refresh` to trigger a rebuild and cursor-based pagination for `graph_changes` and `graph_impact`; `init` registers the server into `~/.claude.json`, `~/.cursor/mcp.json`, `~/.codex/config.toml`, and the Cline CLI settings file (or `.mcp.json` / `.cursor/mcp.json` / `.codex/config.toml` in the repo with `--project`), with `--no-mcp` to omit and remove registrations, `--mcp` to restore, and `--uninstall` to remove only the added entry.

– Flags, files, and agent configs all named precisely. `openqodex@0.11.0`

05 **Report formats, `--report-dir`, and `findings` command** NEW

Adds `--format markdown`, `json`, and `sarif` options to `review`, a `--report-dir <folder>` option to `review` and `scan` for controlling where reports are written (symlinked paths now exit 2 before any write, redirecting all run files away from `.openqodex/`), and a self-contained `report.html` local diff page with findings annotated per line and no external scripts or network calls. Adds `openqodex findings 1,3` to print named findings from the last review in full, and `review` now prints a short receipt by default (verdict, summary, one line per finding with severity/category/file/line) plus absolute paths to `report.html` and `report.md`. Also adds `reviewer.json` (written to the report folder) recording whether a reviewer started and which one, `--block-on-severity`, `--instructions <file>`, and `--reviewer-web on|off` flags to `openqodex review`, and a `scan --block-on-severity <severity>` flag overriding the config's `review.block_on_severity`.

– Flags and output files named with concrete behaviour. `openqodex@0.9.0`

`openqodex@0.7.0` `openqodex@0.5.0`

06 **GitHub Action AI review and hardening** NEW

The GitHub Action gained AI code review: the `ANTHROPIC_API_KEY` secret enables a full Claude Code review alongside the scanners, with findings posted to code scanning and a report in the job summary; a `review` input (`auto` default, `off`, or `required`) controls whether it runs and whether its absence fails the job, a `claude-code-version` input pins the exact Claude Code version installed, and `reviewed`, `review-status`, and `reviewer` outputs report what ran and its outcome. `config-from: head`, `fail-on-tool-error: true`, and `block-on-severity` inputs let the pull request's own config be read, fail the job on scanner tool failure, and set a severity gate the pull request's config cannot weaken, with a `status` output set to `tool-failed` when a scanner tool exits 2. The `version` and `claude-code-`

`version` inputs now accept only exact SemVer strings (e.g. `0.6.1`) — tags like `latest`, ranges, paths, and `file:` references fail the step. The Action hardens its reviewer by reading custom instructions from the base branch only, disabling web tools by default, and restricting PATH to system folders plus `git`, `node`, `npx`, `npm`, and `claude` from outside the checkout; when AI review does not complete, the Action still runs the scanners so no scanner finding is hidden.

– Every input, output, and hardening measure named. `openqodex@0.7.0` `openqodex@0.5.0`

07 Reviewer backend selection and web access IMPROVED

75

Adds `--reviewer codex` to `openqodex review` to run the full review with `codex` `exec` inside a read-only, network-isolated sandbox, with `--reviewer auto` selecting Codex automatically when running inside Codex or when it is the only reviewer installed; `--reviewer auto|claude` and `--timeout <seconds>` flags were also added. Adds a `reviewer:` key (`auto`, `claude`, `codex`, or `cursor`) in `~/.openqodex/config.yaml` to set the reviewer for every review, and a `reviewer_web` key controlling web search and page fetching for reviewers — introduced off by default, then switched to on by default with `reviewer_web: off` to disable it. A `reviewed_by` field in `report.json` and a `run` property in `report.sarif` identify which reviewer completed the review.

– Flags and config keys named, default change tracked. `openqodex@0.6.0` `openqodex@0.5.0`

08 Branch and pull request review targets NEW

75

Adds `openqodex review <branch>` and `openqodex review '<PR link>'` to review a branch or pull request outside your current work — fetched and checked out in a temporary folder, using your own settings and approvals, never the target's — plus `review --finalize --run <id>` to finalize a specific run by name as printed in the brief. Changes to a scanner's own settings or ignore files (such as `.gitignore` or `ruff.toml`) are now raised as reviewer-clearable candidates shown as a note that never counts toward the verdict, and findings are enabled on deletion-only changes: lines adjacent to a deletion now count as changed, with the brief listing each deletion point.

– Exact commands given, mechanism described concretely. `openqodex@0.4.0`

09 `'init'` command overhaul

BREAKING

75

`init` now creates `.openqodex/config.yaml` and `.openqodex/custom-instructions.md` for team-wide conventions, adds a git pre-push hook, and injects a subagent review instruction into each agent's instruction file; config moved to `.openqodex/config.yaml` (created on first run with all keys and defaults) with new keys `review.severity_threshold` (default `minor`) and `review.default_base`, and `pr_review` accepted as an alias for `review`. `--no-repo` skips injecting the review section into `CLAUDE.md` and `AGENTS.md` (`--uninstall` removes exactly

that section), and `init` writes Claude Code rules so `review --agent`, `review --finalize`, their `--all` and `--offline` forms, and `guide` run unattended without confirmation. Later, `--hook none` suppresses the git pre-push hook (with `--no-repo` narrowed to suppressing just the team review section), `init` follows `CLAUDE_CONFIG_DIR` and `CODEX_HOME` environment variables to locate Claude Code and Codex config directories, consolidates its prompts to a single 'Write these files?' question after showing a full plan, auto-detects when running inside Cursor, Claude Code, or Codex CLI agents to skip the terminal requirement, and checks all reviewers simultaneously after writing, reporting `Reviewer ready` or `No reviewer can start yet` with per-reviewer reasons; `init --project` writes agent files inside the repository instead of global scope, and `init --dry-run` prints which scanners would be downloaded and why without downloading. The global instruction line `init` writes to agent files (e.g. `~/.claude/CLAUDE.md`) was replaced with a single sentence, removing the longer section previously written. `init` now also ends with a review of the change, or prompts for what to review, with `--no-review` to skip it, and `init / update` remove receipt records older than 30 days from `~/.openqodex/receipts/`.

– Many flags and env vars named across four releases. openqodex@0.9.0 openqodex@0.5.0
openqodex@0.3.0 openqodex@0.2.0

10 Auto-update engine IMPROVED

75

Adds `openqodex update` with `--now`, `--rollback`, `--off`, `--on`, and `--status` subflags, backed by a persistent launcher installed in `~/.openqodex/bin/` for every user-scope install (replacing direct `npx -y openqodex@<version>` calls in Cursor and Cline rules), which reads the active version from the first line of `~/.openqodex/runtime/current`, falling back to the version `init` installed if that line is missing, malformed, or points to a removed copy. The engine silently installs newer releases (at least 24 hours old, npm-provenance-signed by the release workflow) in the background after a review, scan, or push check run through the launcher, at most once a day, respecting `update: off` in `~/.openqodex/config.yaml` or `OPENQODEX_AUTO_UPDATE=0`, and disabling itself under `--offline` and in CI; `update --rollback` turns updates off before making other changes. It was extended so `update --rollback` instead writes `skip_version: <version>` to `~/.openqodex/config.yaml` rather than disabling updates entirely (rolling back to 0.8.1 or earlier, which cannot read `skip_version`, still disables updates), `update --status` lists pending release notices and stale agent files from the last update, post-update output lists every release notice between old and new versions covering changes to data leaving the machine, push-blocking behaviour, and reviewer configuration, and daily auto-update now gates on agent-contract and config-format compatibility, queuing breaking releases until `openqodex update` is run explicitly.

– Subflags, paths, and gating behaviour all specified. openqodex@0.9.0 openqodex@0.3.0

11 Single-command review with completeness checks

BREAKING

`openqodex review` now runs the full review in one command — snapshotting the change, running scanners and the code graph, starting an isolated reviewer, validating the answer, and printing a single report — and considers a review complete only when every scanner candidate was raised or dropped with a reason and every changed range was covered, otherwise the report names what is missing and exits 2. When no reviewer is available, `review` saves unchecked scanner candidates to a file and exits 2 instead of showing raw scanner output as a review. Push hooks now validate only the exact commits being pushed against what the remote holds (so a force push over unreviewed work is not considered covered) and trust only review records in `~/.openqodex/receipts/`, never report files a branch carries under `.openqodex/`; the git pre-push hook no longer scans or prints scanner findings itself.

– Mechanism and breaking push-hook change both explained. `openqodex@0.5.0`

12 Framework route awareness in review briefs NEW 73

The review brief now lists Django and Rails routes reaching changed code — full paths and names, routes left without a handler, rendered templates and views, changed model migrations, and tests that call, request, or name the changed code — and the code graph reads Express, React, Next.js, FastAPI, and Go `net/http` routes, each with full path, handler function, routers and middleware, components and their rendered children, hooks, and tests that request a route or render a component. The framework plugin API advanced to version 3 so Express, React, and Next.js readers share a single file-walk per JavaScript file, and Python dependencies declared via `-r` or `-c` includes, a `requirements/` folder, or pip-tools `.in` files are now read so Django can be discovered in those layouts.

– Named frameworks and route details, but no command surface. `openqodex@0.11.0`

13 Full scan and review trail recording NEW 73

`report.json` now includes a `timings` object recording millisecond durations for `change_ms`, `scan_ms`, `graph_ms`, `brief_ms`, `reviewer_ms`, `finalize_ms`, and `total_ms` per review run (and `change_ms`, `scan_ms`, `total_ms` per scan), with the `RunTimings` type exported from the library. `scan.json` gains `selection` and `trail` fields recording which scanners a change called for and why, and the disposition of every scanner finding (dropped reason or candidate it became), with each scanner entry gaining `stages`, `resolveMs`, and `files` fields. Semgrep's entries in `report.json` and `scan.json` now list the registry rule packs used, each with its sha256 and byte size, enabling reproducibility checks between scans.

– Field names enumerated but no command to run. `openqodex@0.13.0`

14 Path exclusion matching improvements IMPROVED 60

`review.paths.exclude` now applies to both paths of a renamed file — a file renamed out of an excluded folder is reviewed as new, one renamed into an excluded folder is

reviewed as deleted. Glob patterns in `review.paths.exclude`, `review.disabled_rules`, a lens, or a custom scanner's `paths` are now matched in bounded time, and config globs longer than 1,024 characters are refused with an error naming the file and key.

– Names exact config keys and the 1024-char limit. openqodex@0.13.0 openqodex@0.12.0

15 Reviewer timeout limits raised IMPROVED 60

Raises the default reviewer timeout to 30 minutes for a change, branch, or pull request and 2 hours for `--all` (up from 10 minutes for both); a timed-out review now reports which cap it hit and instructs the user to pass `--timeout <seconds>` to raise it, and `-timeout` values above 2147483 seconds are now refused.

– Concrete numbers and the flag to override them. openqodex@0.13.0

THINNER COVERAGE BELOW

16 Code graph plugin interface bumped to v4 BREAKING 55

The code graph plugin interface is now version 4 and the Express plugin is now version 7; the first build after upgrading will parse every file again rather than using kept facts in `.openqodex/graph/`. Scanner candidate numbering also changed — candidates with the same severity are now ordered deterministically by severity, scanner, file, line, rule, and message — so stored candidate IDs from previous runs will not match the new IDs.

– Explains before/after but no reader action beyond awareness. openqodex@0.12.0

17 Convention file binding for reviews NEW 55

Adds `review.instructions_files` config key so repository owners can specify convention files that are quoted in review briefs, recorded for presence or absence, and bound to `finalize`; the default reads no convention files.

– Single config key named, minimal mechanism beyond that. openqodex@0.13.0

18 `guide` subcommands for agents and privacy NEW 55

Adds `openqodex guide skill` printing the review procedure of the running version (with commands written for the launcher when it started it), `openqodex guide privacy` printing the privacy policy and listing every network call the tool makes (backed by `docs/privacy.md` documenting that nothing is collected), and `openqodex guide for-agents` printing an offline docs page on installing OpenQodex, configuring a repository, adding a custom scanner, and acting on each message.

– Three commands named, each only one line of description. openqodex@0.13.0

openqodex@0.6.1 openqodex@0.3.0

19 Issue reporter command NEW 50

Adds `openqodex report "<what went wrong>"` to interactively draft and optionally create a GitHub issue when the tool or a scanner fails, without sending code, paths, filenames, or secrets.

– One command named, privacy behaviour only briefly described. `openqodex@0.2.0`

20 Codex and Cursor plugin packaging NEW 45

Adds a Codex plugin in `plugins/codex/` carrying the skill for the OpenAI plugin directory and `.cursor-plugin/plugin.json` making the repository a Cursor plugin, and expands the skill's declared capabilities to name the requests it answers: code review, security scan, diff, and pull request.

– File paths named but purpose described thinly. `openqodex@0.6.1`

21 Reviewer usage tracking NEW 40

Adds `usage.json` written to the report folder by `review`, recording each reviewer round's model, token counts, and cost.

– Thin: names the file, no detail on format. `openqodex@0.13.0`

22 Go runtime probe relocated out of home folder IMPROVED 40

The Go runtime probe now runs in `~/.openqodex/cache/runtime-probe` (or `$OPENQODEX_HOME/cache/runtime-probe`) instead of the home folder, suppressing Go telemetry file writes for that probe.

– Minor path change, thinly described. `openqodex@0.13.0`

23 Multi-line secret redaction IMPROVED 30

Multi-line secrets such as private keys are now fully redacted line-by-line wherever each line appears in reports.

– One-line description, no mechanism beyond 'line-by-line'. `openqodex@0.9.0`

↳ BREAKING ON UPGRADE

- ! The code graph plugin interface is now version 4 and the Express plugin is now version 7; the first build after upgrading will parse every file again rather than using kept facts in `.openqodex/graph/`.
- ! The numbering of scanner candidates with the same severity changes in this release (candidates are now ordered deterministically by severity, scanner, file, line, rule, and message), so any stored candidate IDs from previous runs will not match the new IDs.
- ! `graph.max_files` now counts only new parses rather than all files (including those whose facts are cached), which changes its effective meaning for any configuration that set this value expecting it to cap total file reads.
- ! `openqodex update --rollback` no longer disables future updates; it writes `skip_version: <version>` to `~/.openqodex/config.yaml`. Rolling back to 0.8.1 or earlier (which cannot

read `skip_version`) still disables updates.

- ! `doctor --install` now installs only the scanners a repository's files require rather than all scanners; use `doctor --install --all-scanners` to restore the previous behaviour.
- ! brakeman now requires Ruby 3.0 or newer.
- ! The global instruction line `init` writes to agent files (e.g. `~/.claude/CLAUDE.md`) is replaced with a single sentence on the next `init` run; the longer section previously written is removed.
- ! The `version` and `claude-code-version` Action inputs now accept only exact SemVer release versions (e.g. `0.6.1`); a tag such as `latest` , a version range, a path, or a `file:` package now fails the step.
- ! A wrong `config-from` or `block-on-severity` value in the GitHub Action now fails the step rather than being silently ignored.
- ! A tool failure (exit 2) in the GitHub Action no longer fails the job by default; jobs that relied on that behaviour must add `fail-on-tool-error: true` to preserve it.
- ! The git pre-push hook no longer scans or prints scanner findings; workflows that depended on per-push scanner output from the hook will no longer receive it.
- ! The older `review --agent` and `review --finalize` agent rules are removed from the skill and team section; only `review` and `review --all` are authorised for Claude Code without prompting.
- ! Push hooks no longer accept review records from `.openqodex/` files a branch carries; only records in `~/.openqodex/receipts/` are trusted.
- ! `init` and `update` remove receipt records older than 30 days from `~/.openqodex/receipts/` .

openqodex@0.1.0

NOTES ↗

STABLE

OpenQodex 0.1.0 is a new open-source code review tool that runs inside coding agents, with CLI commands, SARIF/JSON/Markdown output, and IDE integrations.

↳ GET THIS VERSION

```
$ git clone --branch openqodex@0.1.0
https://github.com/openqodex/openqodex.git
# already have the repo? check out this version:
$ git checkout openqodex@0.1.0
```

Adds `openqodex review --agent` to analyse a change, select relevant scanners, and print a review brief for the coding agent.

Adds `openqodex review --finalize` to validate agent findings without a model and write the report as Markdown, JSON, and SARIF.

Adds `openqodex scan` for running scanners standalone in git hooks, pre-commit, and CI.

Adds `openqodex init` to install the skill and push gate into Claude Code, Codex CLI, Cursor, and Cline; `--uninstall` removes them.

Adds `openqodex hook install` to add an optional git pre-push hook.

+7 more

WAS THIS USEFUL?



mitmproxy

Feedback



v12.3.0

CODE >

RANK

An interactive TLS-capable intercepting HTTP proxy for penetration testers and software developers.

mitmproxy 12.3.0 adds an eMudhra CA signing option, a dark theme for mitmweb, and small flow-table triage upgrades in the web UI.

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

01 eMudhra CA signing option NEW 80

Adds a `--emudhra` flag to sign with the eMudhra CA, letting testers intercept TLS traffic from clients that still trust the compromised eMudhra intermediate CA — useful for testing or research against legacy trust stores.

Intercept TLS traffic from a client that still trusts the compromised eMudhra intermediate CA — useful for testing or research against legacy trust stores.

```
$ mitmproxy --emudhra
```

— Names exact flag and includes a runnable command. v12.3.0

THINNER COVERAGE BELOW

02 Flow table badges and icons in mitmweb NEW 45

mitmweb's flow table now shows status and method badges for faster at-a-glance triage, and uses SVG icons for flow table resource types.

— Describes two small UI additions but no exact navigation path. v12.3.0

03 Dark theme in mitmweb NEW 15

mitmweb gains a dark theme, per the release summary.

— Mentioned only in passing with no further detail. v12.3.0

WAS THIS USEFUL?



The ZAP by Checkmarx Core project

ZAP Browser Extension 0.3.0 introduces a Script Viewer that lets users inspect, edit, and replay recorded Zest scripts entirely within the browser.

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

01 Script Viewer for Zest recordings in Browser Extension NEW

90

The ZAP Browser Extension 0.3.0 adds a Script Viewer: a **View Script** button in the extension popup opens the recorded Zest script in its own tab, updating statement-by-statement while recording. Statements can be edited in-browser — including changing locator types (**id** , **cssSelector** , **xpath** , **className** , **name** , **linkText** , **partialLinkText** , **tagName**), reordering, deleting, and adding new statements — with Undo/Redo support. **Run in New Window** and **Run in New Tab** replay the script directly in the browser, with a step-through mode (**Step** / **Next**) for statement-by-statement inspection, and **Load** , **Save** , and **New** controls open scripts from disk, download the current script, or clear and start fresh. It ships in both the full ZAP Browser Extension and the standalone Recorder; users launching it from ZAP must update the **Client Side Integration** add-on.

⋮

Verify a recorded login flow is robust before using it in ZAP — step through each statement one at a time to confirm element locators resolve correctly.

1. Update to version 0.3.0 of the ZAP Browser Extension or standalone Recorder. 2. Record your login flow as normal. 3. Click **View Script** in the extension popup — the Script Viewer tab opens. 4. Stop recording, then select **Run in New Tab** . 5. Tick **Step** and press **Next** to advance one statement at a time, watching the on-page indicator highlight each action.

Fix a brittle link locator in a saved Zest script by switching it to **partialLinkText** when link text includes a variable counter.

1. Click **View Script** in the extension popup, then use **Load** to open the saved .zst file from disk. 2. Hover over the relevant statement row and click its Edit button. 3. Change the locator type to **partialLinkText** and enter the stable portion of the link text. 4. Click **Save** to download the updated script. 5. Click **Run in New Tab** to confirm the replay succeeds.

— Names every control and locator type; UI steps given, not a CLI command.

WAS THIS USEFUL?



Cycode

Feedback



LiteLLM 1.107.0: Stopping a Poisoned Package Before It Installs (2026-10-21)

BLOG ↗

RANK

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

Cycode's 1.107.0 release centers on stopping supply-chain attacks before code runs: a new Workstation Protection agent blocks malicious package installs in real time, alongside extensions to SBOM/AIBOM, container scanning, graph queries, and build monitoring to trace and catch poisoned packages.

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

01 Workstation Protection with cooldown and threat-intel blocking

NEW

76

Cycode Workstation Protection (early access), deployable via MDM, intercepts package installs on developer and CI machines in real time. It applies a configurable cooldown policy to hold back recently published package versions, and separately checks a threat intelligence feed to block known-malicious package versions regardless of the cooldown window.

Enforce a cooldown policy so any package version published within the last few days is held on developer and CI machines — blocking a same-day malicious release like LiteLLM 1.107.0 before it installs.

📌 In the Cycode console, go to Workstation Protection > Policies and set the cooldown period to the desired number of days. Deploy the agent to developer machines and CI runners via your MDM provider.

– Mechanism and UI path given, but still early access with few specifics

LiteLLM 1.107.0: Stopping a Poisoned Pa...

02 Cimon eBPF build monitor for CI payload behavior

IMPROVED

65

Cimon, Cycode's eBPF-based build monitor, watches network calls, file writes, and processes on CI runners to detect supply-chain payload behaviors such as `.pth` file writes into `site-packages`.

– Concrete mechanism and example behavior named, but no setup steps

03 Container scanning for poisoned floating-tag images IMPROVED

50

Container scanning now detects poisoned proxy images on floating tags such as `latest`, `main-latest`, and `main-stable` in connected registries.

– Named tags give specificity but no steps to invoke it

LiteLLM 1.107.0: Stopping a Poisoned Pa...

04 Context Intelligence Graph package lookup IMPROVED

50

The Context Intelligence Graph links a finding to its repo, image, pipeline, and team, enabling a single query to list all locations where a specific package version appears and identify its owners.

– Describes capability but no query syntax or UI path given

LiteLLM 1.107.0: Stopping a Poisoned Pa...

05 AIBOM coverage for AI/LLM dependencies IMPROVED

47

SBOM and AIBOM generation now extends to AI/LLM dependencies, surfacing transitive pulls like LiteLLM introduced via agent frameworks such as DSPy and CrewAI.

– Names frameworks and example package but no config or command

LiteLLM 1.107.0: Stopping a Poisoned Pa...

WAS THIS USEFUL?



Endor Labs

Feedback



Endor Labs + SpaceXSI: Securing every stage of agentic software delivery | Endor Labs (2

BLOG

RANK

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

Endor Labs shipped a broad integration with SpaceXSI's Grok Build, adding an in-chat Agent Kit, deterministic governance hooks for agent actions, trust scoring for MCP servers and skills, package/code scanning via MCP and hooks, Package Firewall protection, and AI SAST/SCA scanning for agent-written code.

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

01 Agent Kit for Grok Build

NEW

87

The Endor Labs Agent Kit can be installed via the marketplace or by typing `/add-plugin endorlabs` in agent chat, enabling in-session scanning, triage, and remediation through specialized remediation agents — including an AI SAST Remediation agent and an SCA Remediation agent — that can open pull requests for approved fixes. Agent Kit workflows are also supported outside the IDE, in CI or cloud runs, via the Grok Build SDK.

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

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

– Names exact plugin command, SDK, and two remediation agents

Endor Labs + SpaceXSI: Securing every s...

02 AI SAST and SCA scanning in Grok Build

NEW

70

Endor Labs AI SAST is integrated into Grok Build to scan first-party code for security flaws, trace data flow from untrusted input to vulnerable calls, and flag exposed secrets before commit. Endor Labs SCA is also integrated to check direct and transitive dependencies with reachability analysis and risk scoring, cutting SCA findings by up to 92%.

– Concrete mechanism and 92% figure but no command shown

Endor Labs + SpaceXSI: Securing every s...

THINNER COVERAGE BELOW

03 MCP and hooks integrations for in-session scanning

NEW

An MCP server integration checks for code security risks, exposed secrets, and vulnerable dependencies from within an agent session, and a separate hooks integration scans every package an agent installs, blocking confirmed malicious packages at install time.

– Describes two mechanisms but no endpoint or config named

Endor Labs + SpaceXSI: Securing every s...

04 Coding Agent Governance hooks NEW

50

Coding Agent Governance hooks record every MCP tool call, shell command, and file operation an agent performs and enforce security policies deterministically before each action runs, without depending on a model's judgment.

– Explains mechanism but no config key or command given

Endor Labs + SpaceXSI: Securing every s...

05 Endor Scores for MCP servers and skills NEW

40

Endor Scores are now generated for each MCP server and skill, mapped to the OWASP Top 10 for LLM Applications, giving security teams a trust rating for approved and unapproved agent tooling.

– Names OWASP mapping but no way to view or query scores

Endor Labs + SpaceXSI: Securing every s...

06 Package Firewall extended to Grok Build IMPROVED

38

Endor Labs Package Firewall protection now extends to Grok Build, blocking malicious packages at the registry level for teams that route installs through it.

– Names the product but no setup steps given

Endor Labs + SpaceXSI: Securing every s...

WAS THIS USEFUL?



Cloudsmith

[Feedback](#)[Set exact usage limits with numerical inputs \(2026-10-21\)](#)[BLOG](#)[RANK](#)

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

Cloudsmith replaced slider-based usage limit controls with exact numerical inputs in workspace settings.

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

01 Numerical inputs for workspace usage limits

[IMPROVED](#)

58

In the workspace Settings tab under **Usage limits**, admins can now type exact numerical values for usage limits instead of using the former slider controls, avoiding rounding errors.

Set a precise usage limit for your workspace without rounding errors from slider controls.

 In the Cloudsmith console, go to Workspace Settings > Usage limits and type an exact numerical value into the limit field.

– Names exact UI path but no mechanism beyond input change

[Set exact usage limits with numerical i...](#)

WAS THIS USEFUL?



RANK



Manifest

Feedback



announced · Oct 21, 2026

BLOG

RANK

Manifest secures the software and AI supply chains behind the technology enterprises build and buy, with inside-out visibility across dependencies, binaries, vulnerabilities, provenance, and AI models.

Manifest brought its Foreign Risk contributor-screening capability to self-hosted and air-gapped environments, backed by a local 20-billion-record dataset and audit-ready evidence export.

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

01 Foreign Risk for air-gapped and self-hosted deployments NEW

74

Manifest Foreign Risk now runs entirely inside disconnected environments, running all contributor-level scans against a local copy of a 20-billion-record dataset across 18 risk categories with no data leaving the network. It detects sanctioned-jurisdiction contributors, foreign military/intelligence/state-lab affiliations, and unverifiable contributor identities, and supports exporting findings as audit evidence mapped to NIST 800-53 SR-3 and SR-4 and Cybersecurity Maturity Model Certification (CMMC) supply chain controls. Dataset updates are packaged for local infrastructure and loaded via a repeatable refresh process to keep on-prem findings aligned with the SaaS version without any outbound connection.

– Describes mechanism and named controls but no UI path or command

launch-20261021-60bc1bd8

WAS THIS USEFUL?

infisical

Feedback



v0.167.0

NOTES ↗

RANK

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

Infisical v0.167.0 adds DigiCert HSM-backed code signing support alongside smaller usability improvements to the certificate manager and secrets overview.

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

01 DigiCert HSM key as code-signing signer NEW 40

Supports using an existing DigiCert certificate and its HSM key as a signer in the code-signing workflow.

– Names mechanism but no steps or config shown v0.167.0

02 Copy ID on certificate policy rows NEW 35

Adds Copy ID to alert and approval policy row menus in the certificate manager.

– Clear UI location but minor scope v0.167.0

03 Last-edited timestamps on secrets IMPROVED 30

Shows last-edited time for secrets in the single-environment overview.

– Simple UI addition, minimal detail v0.167.0

WAS THIS USEFUL?



GitGuardian SaaS

Feedback



Catch secrets leaked in public Docker Hub images (2026-10-21)

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 extends Public Monitoring to scan Docker Hub images for leaked secrets, enriches Wiz findings with GitGuardian incident IDs, and adds Windows support for Honeytoken planting.

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

01 Wiz findings enriched with GitGuardian IDs

IMPROVED

73

Findings pushed to Wiz are now enriched with GitGuardian incident, occurrence, and source IDs in the finding's `originalObject` field, enabling use in Wiz workflows and via the Wiz API.

– Names exact field but no further usage steps given

Catch secrets leaked in public Docker H...

02 Docker Hub image scanning in Public Monitoring

NEW

67

GitGuardian Public Monitoring now scans newly published Docker Hub images daily, matching secrets found to your perimeter. Docker Hub incidents surface image-specific context — publishing account, image and repository name, and the layer the secret was found in — instead of commit and author, and are routed through existing notifications, integrations, and agents analysis, matching the workflow already used for GitHub incidents.

– Explains mechanism and scope but no config/flag to act on

Catch secrets leaked in public Docker H...

THINNER COVERAGE BELOW

03 Honeytoken planting on Windows endpoints

IMPROVED

25

Honeytoken planting support is extended to Windows endpoints.

– Bare statement with no mechanism or detail

Catch secrets leaked in public Docker H...

WAS THIS USEFUL?



REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

APKiD

[Feedback](#)[v3.2.0](#)[NOTES ↗](#)[CODE ↗](#)

Android Application Identifier for Packers, Protectors, Obfuscators and Oddities - PEiD for Android

Adds `-st / --scan-trackers`, `-m / --print-matches`, and `--print-filter` CLI flags for tracker scanning and match output filtering.

↳ GET THIS VERSION

```
$ git clone --branch v3.2.0 https://github.com/rednaga/APKiD.git
# already have the repo? check out this version:
$ git checkout v3.2.0
```

NEW DETECTION Adds `-st / --scan-trackers` flag to enable scanning APKs for embedded third-party trackers (generated from the Exodus Privacy tracker database).
Adds `-m / --print-matches` flag to print matched YARA rules and strings to stdout.
Adds `--print-filter RULES` option to filter printed matches by rule name: comma-separated names to include (with `+` or no prefix) or exclude (with `-` prefix); implies `--print-matches`.

UPDATED DETECTION [Google Pairip / Play Integrity rules](#)

UPDATED DETECTION [Protectt YARA rule](#)

NEW DETECTION [Dart >= 3.13 AOT export symbols rule](#)

RETIRED DETECTION [dexguard_d rule removed](#)

UPDATED DETECTION [Appdome rule improvement](#)

↳ BREAKING ON UPGRADE

! The `dexguard_d` YARA rule has been removed; any workflow or downstream tooling that matches on `dexguard_d` results will no longer receive that output.

WAS THIS USEFUL?



NETWORK SECURITY

◆ Email Security

Sublime Security

[Feedback](#)

Platform v1.90.101 (Prompt Injection Detection - Public Beta) (2026-10-21)

BLOG ↗

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

New Prompt Injection Detection (public beta) uses an ML model and SQL function [beta.prompt_injection_detection\(\)](#) to detect and block hidden AI-agent instructions in inbound email, with ASA-confirmed attempts tagged **Email prompt injection** and available for review in the Message List View. Requires Enterprise plan with ASA.

NEW DETECTION

Adds [beta.prompt_injection_detection\(\)](#) SQL function callable from detection rules to score inbound email for prompt injection attempts.

Extends the existing **Send suspicious messages to ASA** automation to route prompt injection detection matches to the Autonomous Security Analyst (ASA) for confirmation and remediation.

Applies the **Email prompt injection** tag to ASA-confirmed prompt injection attempts, enabling filtering in the Message List View by that tag.

WAS THIS USEFUL?



Rspamd

[Feedback](#)[4.2.3](#)[NOTES](#)[RANK](#)

Rapid spam filtering system.

Rspamd 4.2.3 reworks RBL rule evaluation with deduplicated, priority-ranked checks and overhauls the force_actions and Lua pre-result APIs, alongside a handful of breaking fixes to DKIM-domain matching, rule execution, and the Phishtank provider, plus browser-accurate HTML parsing and a new unix_time metadata field.

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

01 Merged RBL checks in rspamd configuration

BREAKING

95

Rspamd 4.2.3 adds `default_merge_checks = true` to RBL config, deduplicating DNS queries when multiple checks resolve the same name and merging their labels into one comma-separated symbol option (e.g. `example.com:dkim,url`) instead of a separate option per check. A per-rule `merge_checks = false` override in `local.d/rbl.conf` restores the old one-symbol-per-check behaviour for specific RBL rules; because non-`one_shot` symbols no longer accumulate once per check, custom RBL rules with multiple checks will score lower when more than one check matches the same IP or domain, and any log parsers, dashboards, or Lua code reading `task:get_symbol()` option strings must be updated to split on commas.

Opt a specific RBL rule out of check merging so it retains per-check symbol insertions and the old scoring behaviour.

```
hcl

rbls {
  my_rbl {
    merge_checks = false;
  }
}
```

Identify which RBL rules have multiple checks before upgrading, to assess scoring impact from the new merge behaviour.

```
$ rspamadm configdump rbl | grep -A5 'checks'
```

Preserve per-check RBL symbol scoring after upgrading, when your composite rules depend on seeing one insertion per matched RBL entry.

```

rbl {
  # ... existing RBL config ...
  merge_checks = false;
}

```

– Full mechanism, config keys, file path and migration commands given. 4.2.3

02 Unix timestamp field in metadata exporter NEW 80

Adds a `unix_time` field to general metadata in `metadata_exporter`, letting exported message metadata include unix timestamps for easier downstream log correlation; configured via `unix_time = true` under `general` in `/etc/rspamd/local.d/metadata_exporter.conf`.

Include unix timestamps in exported message metadata for easier downstream log correlation.

```

metadata_exporter {
  # ... existing exporter config ...
  general {
    unix_time = true;
  }
}

```

– Exact config key, file path and runnable example given. 4.2.3

03 Priority support for force_actions rules NEW 75

`force_actions` rules can now set a `priority` field accepting `low`, `normal`, `high`, and `critical` values to control which rule wins when multiple rules match, with the rank exposed as a `priority:N` symbol option; rules whose `force_actions` depend on composites are now promoted to postfilters as part of this change.

– Named config field and values, but no runnable example. 4.2.3

04 Lua pre-result API priority and return value changes BREAKING 70

Table-form `task:set_pre_result({...})` now defaults to normal priority instead of low and applies all combined flags (e.g. `flags = 'least,no_smtp_message'`) rather than only the first; callers that relied on the old low-priority default must explicitly add `priority = 0`. `task:has_pre_result()` now returns the custom action's own name as its second value instead of the literal string `custom`, so code comparing against `custom` must be updated; it now also returns the action's flags alongside its name.

05 DKIM domain matching fix in RBL checks

BREAKING

65

The `dkim_match_from` flag under `dkim_domainonly` in RBL config was inverted and is now corrected: with `dkim_domainonly = false`, RBL checks now require an exact full `d=` match against the From domain instead of the previous eSLD comparison, while `dkim_domainonly = true` now matches any `d=` sharing the same eSLD. This reverses prior behaviour and may change which DKIM-domain RBL matches fire.

– Behaviour change described precisely but no config example. 4.2.3

06 RBL rule-matching and provider bug fixes

BREAKING

65

`require_symbols` rules that previously never ran now execute and insert symbols whenever all listed symbols are present; `return_codes` / `return_bits` Lua patterns that were previously compared literally (and never matched) are now matched as patterns; and `phishtank.rspamd.com` is now disabled by default, so setups relying on Phishtank lookups will silently lose that check unless re-enabled.

– Three named fixes described but no reconfiguration example. 4.2.3

THINNER COVERAGE BELOW

07 Browser-accurate HTML parsing

IMPROVED

15

Introduces browser-accurate HTML parsing for more faithful rendering fidelity in spam analysis.

– Bare description with no mechanism or usage detail. product docs

↳ **BREAKING ON UPGRADE**

- ! `default_merge_checks = true` is now the default in RBL config: custom RBL rules with multiple checks whose symbols are not `one_shot` (e.g. `checks = ["from", "received"]`) will score lower when more than one check matches the same IP or domain.
- ! RBL symbol options now emit comma-separated labels in a single option (e.g. `example.com:dkim,url`) instead of separate options per check — any log parsers, dashboards, or Lua code reading `task:get_symbol()` option strings must be updated to split on commas.
- ! `require_symbols` rules that previously never ran will now execute and insert symbols whenever all listed symbols are present.
- ! `return_codes` / `return_bits` Lua patterns that were previously compared literally (and never matched) are now matched as patterns.
- ! `dkim_match_from` with `dkim_domainonly = false` now requires an exact full `d=` match against the From domain (querying the full `d=`), reversing the previous eSLD comparison behaviour; `dkim_domainonly = true` now matches any `d=` sharing the same eSLD.

- ! Table-form task: `set_pre_result({...})` now defaults to normal priority instead of low — callers that relied on low priority must explicitly add `priority = 0`.
- ! task: `has_pre_result()` now returns the custom action's own name as its second value instead of the string `custom` — code comparing against `custom` must be updated.
- ! `phishtank.rspamd.com` is now disabled by default; setups relying on Phishtank lookups will silently lose that check unless re-enabled.
- ! RBL duplicate DNS queries are now merged into a single symbol insertion, so a rule scores only once per matched value instead of once per check. Set `merge_checks = false` to restore the previous per-check scoring behaviour.
- ! RBL: the `dkim_domainonly` flag under `dkim_match_from` was inverted; eSLDs were compared when `dkim_domainonly` was off. The logic is now corrected, which may change which DKIM-domain RBL matches fire.

WAS THIS USEFUL?



Cisco Snort3

[Feedback](#)

3.12.4.0

[NOTES ↗](#)[RANK](#)

Snort++

Snort3 3.12.4.0 adds brotli decompression and `RST_STREAM` rate limiting to HTTP inspection, expands `stream_tcp` sequence tracing, and introduces JSON stats export alongside several smaller observability and logging additions.

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

01 TCP sequence tracing and normalization in `stream_tcp` NEW

70

Adds a `trace_seq_relative` option to `stream_tcp` to toggle between relative and absolute TCP sequence numbers in traces, and adds asymmetric-specific sequence span normalization checks to `stream_tcp`.

Use relative sequence numbers in TCP traces to make `stream_tcp` output easier to read during a session investigation.

```
lua

stream_tcp = { trace_seq_relative = true }
```

– Exact config key with a runnable example, but no mechanism depth. [3.12.4.0](#)

02 Dynamic enable/disable for buffered logging in `logger` NEW

65

Adds dynamic enable/disable commands for buffered logging with configurable `stop_after_match` to `logger`.

Enable `stop_after_match` on buffered logging to halt log output after the first matching event, reducing noise in high-volume environments.

```
lua

logger = { stop_after_match = true }
```

– Has exact config key and a runnable example. [3.12.4.0](#)

THINNER COVERAGE BELOW

03 Stats and peg counter additions across modules NEW

55

Adds support for dumping stats in JSON format via `stats`, adds `peg counters` to `file_api` for extended file observability, and adds a `peg counter` for max header count to `http_inspect`.

– Named counters and output format but no example or schema. 3.12.4.0

04 **RST_STREAM rate limiting in http2_inspect** NEW 50

Adds a configurable `RST_STREAM` rate threshold to `http2_inspect` to cap `RST_STREAM` abuse.

– Names the module and purpose but not the exact option key. 3.12.4.0

05 **Multi-destination PacketTracer output helpers** NEW 50

Adds `stream_log` and `flow_log` helpers to `stream` for multi-destination PacketTracer output.

– Names both helpers but no usage example given. 3.12.4.0

06 **Configurable response pages per action in react** NEW 40

Supports configurable response pages per action variable in `react`.

– Describes capability but omits exact config syntax. 3.12.4.0

07 **Brotli decompression support in http_inspect** NEW 35

Adds `brotli` decompression support to `http_inspect`.

– Single bare line with no usage or configuration detail. 3.12.4.0

08 **QinQ outer VLAN tag support** NEW 35

Adds QinQ outer VLAN tag support to `main`.

– Bare one-line addition with no further detail. 3.12.4.0

WAS THIS USEFUL?  

Sophos DNS Protection

[Feedback](#)

2026.44Released on October 21, 2026

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

Sophos DNS Protection is a cloud-based service that filters DNS requests to block access to malicious, phishing, and unwanted websites across a network.

Sophos DNS Protection shipped automated Firewall onboarding for SFOS 23.0 and new DNS activity widgets on the Sophos Fusion global dashboard.

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

01 DNS Protection widgets on Fusion dashboard NEW 80

Adds DNS Protection widgets to the Sophos Fusion global dashboard — including **DNS Protection summary**, **Top domain queries**, and **DNS usage from protected computers** — reachable via **Dashboards > Manage dashboards > Create new** using the **Product** filter (select **DNS Protection**) or **Category** filter (select **Threat**).

Monitor DNS activity across your environment from the Sophos Fusion global dashboard by adding the DNS Protection summary widget.

1. Go to Dashboards > Manage dashboards and click Create new. 2. In the Product filter, select **DNS Protection** (or in the Category filter, select **Threat**). 3. Add the **DNS Protection summary**, **Top domain queries**, or **DNS usage from protected computers** widget to your dashboard.

— Names exact widgets and navigation path to enable them

2026.44Released on October 21, 2026

THINNER COVERAGE BELOW

02 Automated Sophos Firewall onboarding in SFOS 23.0 NEW 50

Automates Sophos Firewall onboarding and DNS over HTTPS (DoH) configuration in SFOS 23.0, eliminating manual location creation in Sophos Fusion and manual DNS Protection IP address configuration on the firewall.

— Explains before/after but no exact steps or config keys

2026.44Released on October 21, 2026

WAS THIS USEFUL?



CLOUD & CONTAINER SECURITY

Plerion

[Feedback](#)[changelog · Oct 22, 2026](#)[CHANGELOG >](#)[RANK](#)

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

Plerion's latest release adds dedicated Lambda Web Function security coverage — four new checks and a new inventory view — alongside a deeper Sumo Logic workflow integration, a new APEC CBPR compliance framework, a standalone Agent Registry service, and an expanded Global CBPR mapping.

←→ **WHAT SHIPPED · 6 FEATURES** most completely described first [? what's the number?](#)

01 Sumo Logic workflow integration and audit log streaming

 90

BREAKING

Adds a Sumo Logic workflow integration that posts findings as they open, change, and resolve (with a weekly full resend), supports **Send tenant audit logs** to stream audit events in real time, and provides a **Sync now** option for immediate delivery. A new **Add audit log condition** workflow condition lets tenant audit logs be sent independently or alongside finding conditions, replacing the former **Send tenant audit logs** checkbox on the Sumo Logic action; existing workflows that had the box ticked continue sending audit logs automatically.

Stream tenant audit logs to Sumo Logic in real time alongside findings, using the new workflow condition instead of the retired checkbox.

📌 In the Plerion console, open Workflows, select or create a workflow with a Sumo Logic action, click **Add audit log condition**, then save — audit events will post to your Sumo Logic HTTP Source as they occur.

– Names config options and gives migration note with UI steps [changelog-20261022-10b6f4ce](#)

02 Four new Lambda Web Function security checks **NEW**

 85

Adds four checks for Lambda Web Functions: **PLERION-AWS-1246** (High) flags **ApplicationManaged** authentication-type endpoints accepting unauthenticated internet requests, naming the endpoints and domains and surfacing them in exposure views and attack paths; **PLERION-AWS-1247** flags resource policies granting external AWS principals outside your Organization invoke or management access, rating High for unrecognized accounts and Medium for trusted or conditional grants, and showing principal, actions, and granting statement in Access Grants; **PLERION-AWS-1248** (Medium) flags execution roles that allow every action on every resource with no condition, naming the role and policy statements; and **PLERION-AWS-1249** (Low) flags

execution roles with wildcard service-level permissions such as `s3:*` and no condition, naming the role and policy statements.

– Names all four check IDs, severities and logic, but no remediation command

`changelog-20261022-10b6f4ce`

03 APEC CBPR compliance framework NEW

85

Adds APEC Cross-Border Privacy Rules (CBPR) System Program Requirements (November 2019) as a new compliance framework, mapping 6 of 73 requirements covering security safeguards, attack detection/response, and safeguard testing across AWS, Azure, GCP, and Kubernetes; enable it from Compliance under **Manage frameworks**.

Enable the APEC CBPR compliance framework to start mapping your cloud posture against its security safeguard requirements across AWS, Azure, GCP, and Kubernetes.

 In the Plerion console, go to **Manage > Compliance > Manage frameworks**, locate **APEC Cross-Border Privacy Rules (CBPR) System Program Requirements**, and toggle it on.

– Exact coverage numbers and exact enablement path given `changelog-20261022-10b6f4ce`

04 Agent Registry service migration IMPROVED

75

Migrates check `PLERION-AWS-1110` ('Agent registry publishes new records without review') to a new standalone Agent Registry service, collecting registries via `agent-registry` read permissions added in the latest Plerion stack, and flags registries whose auto-approval rules include `APPROVE_ALL`.

– Names check ID, permission and rule but no UI steps `changelog-20261022-10b6f4ce`

05 Global CBPR compliance mapping refresh IMPROVED

60

Refreshes the Global CBPR (September 2025 edition) mapping, expanding coverage across AWS, Azure, GCP, and Kubernetes for its 7 mapped requirements with checks for backups, encryption, exposure, audit logging, vulnerabilities, secrets, and alerting.

– Gives requirement count and coverage areas but no enablement steps

`changelog-20261022-10b6f4ce`

THINNER COVERAGE BELOW

06 Lambda Web Functions asset inventory NEW

50

Adds Lambda Web Functions to asset inventory, showing each function's endpoints with authentication type and domain, revisions, execution roles, resource policy, and tags.

– Lists inventory fields but gives no UI path or command `changelog-20261022-10b6f4ce`

! The **Send tenant audit logs** checkbox on the Sumo Logic action is replaced by the **Add audit log condition** workflow condition; existing workflows that had the box ticked continue sending audit logs automatically.

WAS THIS USEFUL?



Upwind

Feedback



Upwind Story Replay: Watch the Attack Unfold (2026-10-21)

BLOG ↗

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

Upwind Story Replay delivers a visual, MITRE-annotated attack timeline so analysts and CISOs can follow multi-stage threats from first detection to remediation at a glance.

Adds Upwind Story Replay, a visual timeline inside the Upwind console that renders every detection contributing to a threat in chronological order, annotated with MITRE tactics and associated resources.

Story Replay surfaces the order, timing, and blast radius of multi-stage attacks across workloads and identities, replacing manual alert-list reconstruction with a single visual attack path.

NEW VULNERABILITY CHECK

PHP core unauthenticated RCE detection rules (CVE-2026-109618, CVE-2026-109619)

NEW DETECTION

eMudhra intermediate CA key exposure detection

NEW DETECTION

LiteLLM 1.107.0 poisoned PyPI backdoor detection

NEW VULNERABILITY CHECK

GitLab integer overflow RCE behavioral detections (CVE-2026-110264)

WAS THIS USEFUL?



Buildah

[Feedback](#)

v1.46.0

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

A tool that facilitates building OCI images.

Buildah v1.46.0 adds several new flags for fine-grained control over `buildah run`, `add` and `copy`, improves Git source handling and cache mount sharing, and surfaces CDI device info in `buildah info`.

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

01 `--keep-git-dir` preserves `.git` on add NEW 80

Implements `--keep-git-dir` flag on `buildah add` (ADD instruction) for Docker-compatible preservation of the `.git` directory when adding a Git source.

Add a Git repo to a container image while preserving the `.git` directory, enabling version-aware tooling inside the image.

```
$ buildah add --keep-git-dir mycontainer  
https://github.com/example/repo.git /src
```

– Named flag, exact command, clear behavior stated. v1.46.0

02 File selection flags for `buildah add`/`copy` NEW 75

Adds `--include` flag to `buildah add` and `buildah copy` to filter which files are copied from a source, and `--required-path` flag to `buildah add` and `buildah copy` to require that a specific path exists in the source.

Copy only specific files from a directory into a container, reducing image bloat without writing complex `.dockerignore` rules.

```
$ buildah copy --include '*.conf' mycontainer ./configs  
/etc/myapp/
```

– Both flags named verbatim with a runnable example for one. v1.46.0

03 `--umask` flag for `buildah run` NEW 70

Adds `--umask` flag to `buildah run` to set the umask for the container process, useful when build scripts create files with strict permission requirements.

Control the umask of processes launched in a build step, useful when build scripts create files with strict permission requirements.

```
$ buildah run --umask 0022 mycontainer -- make install
```

– Named flag with runnable example, limited mechanism detail. v1.46.0

04 Git checksum verification and cache hits for `buildah add`

IMPROVED

65

Improves Git source handling in `buildah add` by adding `--checksum` verification against the resolved commit and enabling build cache hits for repeated Git URL adds.

– Names flag and mechanism but lacks a runnable example. v1.46.0

THINNER COVERAGE BELOW

05 `sharing=private` cache mounts in `RUN --mount=type=cache`

NEW

55

Adds `sharing=private` support for cache mounts in `RUN --mount=type=cache`.

– Named config value but no example or deeper mechanism. v1.46.0

06 Silence build output with `--log-level`

IMPROVED

50

Enables `--log-level error` or `--log-level fatal` to silence Containerfile step output during builds.

– Named flag values, no runnable example provided. v1.46.0

07 Proxy settings from `containers.conf` honored

IMPROVED

45

Honors `http_proxy` and proxy environment settings from `containers.conf` during image builds.

– Names config source but gives no usage example. v1.46.0

08 CDI visibility in `buildah info`

NEW

25

Adds CDI visibility to `buildah info`.

– Mentioned only in summary with no further detail. v1.46.0

WAS THIS USEFUL?



Alibaba Cloud OpenSandbox Feedback ⓘ

release-1.2.0-rc.1

NOTES ↗

RANK

Secure, Fast, and Extensible Sandbox runtime for AI agents.

OpenSandbox 1.2.0-rc.1 introduces opt-in auto-pause and wake-on-access for fast sandboxes (with a matching SDK wake-retry contract), an experimental Vault-backed egress sidecar, and execd-level OOM and sandbox-scoped resource metrics, alongside breaking changes to idle-pause server config and BatchSandbox restart behavior.

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

01 Auto-pause and wake-on-access for fast sandboxes NEW

90

The `lifecycle.idlePolicy` create field (`idleTimeoutSeconds` 30–86400, `onIdle: pause` , `wakeOnAccess` , optional `activeProbe`) opts fast sandboxes into server-side auto-pause and wake-on-access (OSEP-0024); it persists as the `sandbox.opensandbox.io/idle-policy` Sandbox CR annotation and is off by default. The `OpenSandbox-Skip-Activity: skip` request header suppresses ingress activity observation for a single request (e.g. health checks), stripped at both the gateway and server-proxy hops. All five SDKs add a wake-retry contract: on a `503` carrying the `OpenSandbox-Wake: expired` gateway header, the SDK resends after `Retry-After` within a bounded budget (default 5 attempts / 60 s, configurable per client); exhaustion raises a typed wake-timeout exception carrying attempt count and waited time.

Create a fast sandbox that auto-pauses after 5 minutes of inactivity and wakes transparently when the next request arrives.

```
json
{
  "lifecycle": {
    "idlePolicy": {
      "idleTimeoutSeconds": 300,
      "onIdle": "pause",
      "wakeOnAccess": true
    }
  }
}
```

Probe an auto-paused sandbox from a health-check script without recording the request as user activity, keeping the idle timer accurate.

```
$ curl -H 'OpenSandbox-Skip-Activity: skip' https://<sandbox-host>/healthz
```

– Exact config fields and headers given; SDK retry mechanism described, not shown in code.

release-1.2.0-rc.1

02 OOM kill reporting in execd NEW 80

execd reports OOM kills via cgroupfs polling: it polls the sandbox cgroup every 5 s, diffs `memory.events` (cgroup v2; best-effort v1), and exports the `execd.system.oom.kills` metric with a `local` / `descendant` level attribute, plus structured warn logs with best-effort victim attribution.

– Metric name and polling mechanism given, no query example.

release-1.2.0-rc.1

03 Sandbox-scoped CPU and memory metrics in execd IMPROVED 80

When the sandbox cgroup carries configured limits, `execd.system.cpu.usage` and `execd.system.memory.usage_bytes` now read from the sandbox's own cgroup (`cpu.stat` / `cpu.max`, `memory.current` / `memory.max`, with v1 fallback) instead of the host-wide `/proc` view via gopsutil.

– Exact metric names and source change given, no usage example.

release-1.2.0-rc.1

04 Idle-pause server config section renamed, Redis DSN required 75

BREAKING

Idle-pause configuration moves from `[idle_pause]` to `[automation.idle_pause]`; the redundant `redis.enabled` gate is removed and `redis.dsn` is now required and validated at load when `enabled = true`, turning a previously silent non-start into a loud failure.

Migrate the server config to the new idle-pause section layout and enforce a valid Redis DSN so the sweeper starts loudly rather than silently skipping.

```
ini

[automation.idle_pause]
enabled = true
redis.dsn = "redis://localhost:6379/0"
```

– Exact section path and required key given with migration example.

release-1.2.0-rc.1

05 Pasta, slirp4netns and none network modes NEW 65

`pasta`, `slirp4netns`, and `none` are now valid container-runtime network modes, passed through to podman kernel-level networking; `networkPolicy` is rejected for these modes since no bridge is available for the egress sidecar.

– Named modes and rejected field, no exact config key path shown. release-1.2.0-rc.1

06 BatchSandbox restartPolicy defaults to Never BREAKING 65

BatchSandbox pods now default to `restartPolicy: Never`; workloads relying on the previous implicit restart behavior must set an explicit `restartPolicy` in the pod template to preserve it.

– Clear migration note but limited mechanism detail. release-1.2.0-rc.1

07 BatchSandbox retention during teardown IMPROVED 60

A BatchSandbox is retained in the lifecycle API (reporting `Stopping`) until foreground cascading deletion and cleanup finalizers complete, preventing a replacement sandbox from writing to the same persistent directory before the old writer exits.

– Mechanism and benefit explained, no control flag exposed. release-1.2.0-rc.1

THINNER COVERAGE BELOW

08 Experimental egress sidecar with Vault admission NEW 55

Adds an experimental opt-in egress sidecar with live credential-bound Vault admission, HTTP/2 support with bounded drain, atomic policy-file persistence, and runtime fault containment (OSEP-0023).

– Named mechanisms but no opt-in flag or config shown. release-1.2.0-rc.1

! --> BREAKING ON UPGRADE

! The `[idle_pause]` config section is renamed to `[automation.idle_pause]`; the `redis.enabled` gate is removed and `redis.dsn` is now required (validated at load) when `enabled = true` — a previously silent non-start now fails loudly.

! BatchSandbox pods now default to `restartPolicy: Never`; any workload relying on the previous implicit restart behavior must set an explicit `restartPolicy` in the pod template to preserve it.

WAS THIS USEFUL?



DETECTION & RESPONSE

◆ Security Data Pipeline

Cribl

[Feedback](#)[changelog · Oct 22, 2026](#)[DOCS ↗](#)[BLOG ↗](#)

Cribl is the AI Platform for Telemetry, providing open, flexible infrastructure that helps enterprises manage, connect, and analyze their data across IT, security, and observability environments.

Cribl Search gains optional OCSF v1.8 field mapping in the Lakehouse Engine via dataset rules, with modes to add or replace fields and support for custom Datatypes.

NEW COMPLIANCE FRAMEWORK

Adds OCSF v1.8 field mapping to Cribl Search's Lakehouse Engine, configurable via a dataset rule (GDI Step 3), so analysts can query security data using consistent OCSF field names instead of vendor-specific ones like `src_ip`, `source_ip`, or `srcaddr`.

Supports two mapping modes: add mapped OCSF fields alongside original source fields, or replace source fields with OCSF names — configurable per dataset rule.

Allows creation and cloning of custom Datatypes in Cribl Search to extend or override built-in OCSF Schema Mappings for any desired schema.

WAS THIS USEFUL?



Meta ThreatExchange

Feedback



hma-v1.2.1

NOTES ▸

RANK

Trust & Safety tools for working together to fight digital harms.

ThreatExchange's HMA tool added a new CLI query command for server-side filtered ThreatExchange searches and improved fetch status reporting to show exchange enabled state.

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

01 Server-side filtered ThreatExchange query command NEW

75

Adds `threatexchange api fb_threatexchange query` to pytx, allowing a server-side filtered search against ThreatExchange without pulling down the full dataset locally.

Run a server-side filtered search against ThreatExchange without pulling down the full dataset locally.

```
$ threatexchange api fb_threatexchange query
```

– Names exact runnable command, light on mechanism detail hma-v1.2.1

THINNER COVERAGE BELOW

02 Exchange enabled state in HMA fetch status IMPROVED

35

HMA's fetch status now reports the enabled state of exchanges, making it visible whether a given exchange is active.

– Describes what changed but no surface or command named hma-v1.2.1

WAS THIS USEFUL?



OpenCTI

Feedback



7.261021.1

NOTES

RANK

Open Cyber Threat Intelligence Platform

OpenCTI 7.261021.1 adds in-chat downloads of XTM One-generated decks, reports and videos via Ask Ariane, a new configuration export/preset module, and incremental improvements to XTM Hub messaging, security coverage, custom fields and graph surfaces.

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

01 Config-preset export module in Filigran Experience UI NEW 80

A new **config-preset** module adds an export configuration drawer in the Filigran Experience UI that lets users select exportable types and bundle platform configuration into a downloadable ZIP, backed by a bypass-only storage root and an exportable types query.

Export your OpenCTI configuration as a preset bundle to share or back up platform settings via the Filigran Experience UI.

In the OpenCTI UI, go to Filigran Experience > Export Configuration, select the types to include, and click Export — OpenCTI generates and stores a config bundle ZIP you can download.

— Names module, storage mechanism and UI path with a worked example 7.261021.1

THINNER COVERAGE BELOW

02 Ask Ariane deck, report and video downloads NEW 45

Ask Ariane can now view and download decks, reports, and videos generated by XTM One directly inside the chatbot.

— States capability but no mechanism or UI path 7.261021.1

03 Custom field value management in playbooks and mass operations IMPROVED 40

Adds custom field value management in playbooks and mass operations, described as chunk 7 of an ongoing custom-fields rollout.

— Names surfaces but omits specific fields or steps 7.261021.1

04 XTM Hub browser message parameter IMPROVED 25

Adds a parameter to XTM Hub browser messages for additional control over hub interactions.

– Parameter named only generically, not by name 7.261021.1

05 Manual security coverage result updates IMPROVED 20

Enables updating a manual security coverage result.

– Single vague sentence with no mechanism 7.261021.1

06 Graph experience upgrades across all surfaces IMPROVED 15

Elevates the graph experience across every graph surface in the platform.

– No concrete detail on what changed in the graphs 7.261021.1

WAS THIS USEFUL?



Cloudmersive

Feedback



changelog · Oct 22, 2026

CHANGELOG ↗

RANK

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

Cloudmersive expanded document class support in its Private Cloud Fraud AI API and added remote update management to its Private Cloud Installer across its core APIs.

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

01 Remote update management in Private Cloud Installer

NEW

55

Adds remote management support for **Install Updates** in the Private Cloud Installer v3-73 for Virus Scan API, Convert API, Validate API, and NLP API.

- Names flag and version but no mechanism detail

changelog-20261022-4250cd6f

02 Expanded document classes in Fraud AI API

IMPROVED

30

Expands supported document classes in the Private Cloud Fraud AI API (v2-2-0 and v2-3-0).

- No details on which document classes were added

changelog-20261022-4250cd6f

WAS THIS USEFUL?



Sumo Logic

[Feedback](#)[In-App Tabs for Dashboards and Metrics \(Search\) \(2026-10-22\)](#)[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 extended in-app tabs to cover Dashboards and Metrics alongside Log Search, and overhauled retention options for metrics transformation rules with automatic migration of existing 15-day rules to 400 days.

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

01 In-app tabs for Dashboards and Metrics **IMPROVED** 65

In-app tabs in the New UI now support Dashboard and Metrics tabs alongside Log Search, letting all three stay open in a single browser window without losing context. Requires the **Enable in-app tabs** preference to be turned on, found in user Preferences.

Enable in-app tabs so you can keep Log Search, Dashboard, and Metrics open simultaneously without extra browser tabs.

📌 In the New UI, go to your user Preferences and enable the **Enable in-app tabs** option. You can then open Log Search, Dashboard, and Metrics tabs side by side and switch between them without losing context.

– Names the config preference and UI path to enable it.

[In-App Tabs for Dashboards and Metrics ...](#)

02 400-day retention for metrics transformation rules **BREAKING** 65

Metrics transformation rules now offer **Do Not Store** or **400 days** as retention options for matching metrics, replacing the previous 15-day option, and standardize aggregate retention for aggregated metrics to **400 days**. Existing metrics transformation rules configured with 15-day retention are automatically migrated to 400-day retention.

– Clear before/after and migration behavior, but no config key named.

[Updated Retention Options for Metrics T...](#)

↳ **BREAKING ON UPGRADE**

! Existing metrics transformation rules configured with 15-day retention are automatically migrated to 400-day retention.

WAS THIS USEFUL?



Databricks

Feedback



Screen MCP tool calls and results for jailbreaks (Beta) (2026-10-21)

BLOG ↗

RANK

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

Databricks added a Beta jailbreak guardrail to Unity Gateway that screens MCP tool calls and their results for prompt injection attempts.

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

01 Jailbreak guardrail for MCP tool calls in Unity Gateway NEW

45

Unity Gateway now supports a Jailbreak guardrail (Beta) that screens MCP tool calls and their results for jailbreak and prompt injection attempts, configurable to apply on input, output, or both.

– Names the guardrail and scope but no config key or API given.

Screen MCP tool calls and results for j...

WAS THIS USEFUL?



Splunk Security Content

Feedback



v6.9.0

NOTES >

RANK

Splunk Security Content

Splunk Security Content v6.9.0 ships 15 new detections spanning Cisco AI Defense, AWS CloudTrail, Zeek network traffic, and endpoint telemetry, introduces a new GenAI abuse analytic story and an alpha Cisco AI Defense Threat Operations dashboard, and refreshes several existing detections and lookups.

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

01 Cisco AI Defense expansion NEW

80

Adds Cisco AI Defense Allowed Sensitive Prompt , Cisco AI Defense Allowed Sensitive Response , and Cisco AI Defense Repeated High Risk Activity detections along with a new Abuse of GenAI Systems and Applications analytic story, and introduces the Cisco AI Defense Threat Operations dashboard (alpha) for reviewing runtime activity, enforcement outcomes, risky behavior, sensitive data matches, and activity by application and model.

– Names detections and story but no detection logic shown Describes dashboard scope but alpha status lacks exact navigation

v6.9.0

02 New cross-platform detections NEW

70

Adds AWS Anomalous Secrets Manager Batch Retrieval , AWS Repeated CloudWatch Logs Read Operations , AWS EC2 Multiple SSH Public Keys Uploaded by a Single Identity , AWS EC2 Startup Script Modified on Stopped Instance , and AWS IAM Role Created with External Trust AWS CloudTrail detections; HTTP POST Request with Encoded Payload , Unusual SSH Client Banner , and Unusual HTTP Request Method Zeek-based detections with new Zeek HTTP and SSH data sources and macros; and PowerShell Download from Code Hosting Service and Windows Process Outbound Connection To Social Media Domain endpoint detections.

– Five named detections but no rule logic or thresholds given Names detections and data sources but no field or macro detail Bare detection names, no behavioral detail

v6.9.0

THINNER COVERAGE BELOW

03 Braodo Stealer coverage update IMPROVED

40

Adds two new analytics and tags existing detections for the 2026 Braodo Stealer variant under the updated **Braodo Stealer** analytic story.

– Vague on which analytics or detections were tagged v6.9.0

04 Remote access and file-sharing lookups update IMPROVED

 40

Refreshes remote access software coverage and its lookup, adds a file-sharing domains lookup, and updates suspicious-port detection for Cisco Secure Firewall.

– Describes changes but no lookup names or detection IDs v6.9.0

05 atomic_guid updates for Atomic Red Team linkage IMPROVED

 35

Updates **atomic_guid** values across detections to link with relevant Atomic Red Team tests.

– Names the field but not scope of detections affected v6.9.0

06 M365 Copilot jailbreak analytics to production IMPROVED

 25

Advances four M365 Copilot jailbreak analytics to production status.

– No names of the analytics or what changed in status v6.9.0

WAS THIS USEFUL?



Tenzir Node v6.20.0-rc.4: A new data model for faster pipelines (2026-10-21)

BLOG ↗

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

New heterogeneous-list data model boosts pipeline speed; adds `ai_decide`, encryption/checksum functions, `type_kind`, secrets in function args, and parallel load balancing.

Introduces a new data model enabling heterogeneous lists (lists holding values of different types) and faster pipelines.

NEW INTEGRATION Adds `ai_decide` operator with typed decisions and a dedicated **Decisions** endpoints configuration.

Adds `type_kind` function for inspecting value types in pipelines.

Adds encryption functions and support for secrets as function arguments, including **Secret keys for hmac** and Parquet encryption with a shared AES key.

Adds **Luhn check digits** and **CRC-32 checksums** functions.

+12 more

NEW INTEGRATION [Read from and write to DuckDB databases](#)

NEW INTEGRATION [Send events to Hydrolix](#)

NEW INTEGRATION [Read events from Hydrolix](#)

NEW INTEGRATION [Windows Event Forwarding without a Windows collector](#)

NEW INTEGRATION [Write metrics to Google Cloud Monitoring](#)

UPDATED INTEGRATION [SentinelOne Data Lake API keys and summaries](#)

NEW COMPLIANCE FRAMEWORK [OCSF functions](#)

WAS THIS USEFUL?



SentinelOne

Feedback



announced · Oct 21, 2026

PRESS >

RANK

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

SentinelOne launched Singularity ASOC, a unified SOC suite bundling AI SIEM, Agentic Investigation, Hyperautomation and AI Data Pipelines with 50+ third-party connectors, alongside Singularity ADR to discover, trace and govern AI agents at runtime, and extended Prompt Security's AI governance into Lenovo's ThinkShield endpoint portfolio.

←→ WHAT SHIPPED · 14 FEATURES most completely described first ? what's the number?

01 AI Data Pipelines with 50+ connectors NEW 75

Adds AI Data Pipelines with more than 50 out-of-the-box connectors to third-party sources spanning cloud, identity, email, network, SaaS, data platforms, and vulnerability management — including AWS , Google Cloud , Microsoft , Okta , Cisco , Proofpoint , Netskope , Snowflake , Atlassian , Oracle , Tenable , Microsoft Defender XDR , Palo Alto Networks Prisma , Wiz , Darktrace , Sophos Central , and Check Point . Supports ingestion via syslog and Kafka today, with OpenTelemetry and generic API collection planned, and can filter and route data to third-party SIEMs and storage destinations in addition to Singularity Data Lake .

– Names every connector and protocol but no config steps to wire one up

launch-20261021-2547d6a6

02 Real-time AI agent protection NEW 75

Adds in-line prevention that blocks prompt injection, jailbreaks, credential exfiltration, and tool poisoning before the action executes; intent drift detection that flags the gap between what an agent was asked to do and what it actually did; and agent containment that stops a single agent mid-action without halting the rest of the workflow.

Contain a single rogue agent mid-action — for example one exfiltrating credentials — without stopping the broader agentic workflow.

📌 In the Singularity Platform console, go to ADR > Agent Inventory > locate the offending agent > select Contain Agent to stop its actions while leaving sibling agents in the workflow running.

– Names three distinct mechanisms with a working containment path

launch-20261021-7d5f1e18

03 AI agent and MCP inventory across platforms NEW 70

Discovers every AI agent, tool, skill, connector, and MCP server across endpoints, SaaS, and cloud providers — sanctioned or not — along with the permissions each holds and the employees behind it. Coverage spans Salesforce Agentforce, Microsoft Copilot Studio, OpenAI ChatGPT workspace agents, and Anthropic Claude managed agents, going beyond MCP gateway coverage.

– Names platforms covered but no UI path to view the inventory [launch-20261021-7d5f1e18](#)

04 **Unified Storyline for agentic and EDR events** NEW 65

Integrates ADR event behavior with traditional EDR behavior in a single execution graph, showing the user request, AI assistant, MCP server, and resulting file or process in one chain.

Review the Unified Storyline for an AI agent-initiated incident to see the full chain from user prompt through MCP server to process execution in one graph.

 In the Singularity Platform console, go to Incidents > select an ADR-flagged event > open the Storyline view to trace the user request, AI assistant, MCP server, and resulting file or process in a single execution graph.

– Includes a console navigation path to view the chain [launch-20261021-7d5f1e18](#)

THINNER COVERAGE BELOW

05 **Prompt Security AI governance in Lenovo ThinkShield** NEW 50

Adds SentinelOne Prompt Security to the Lenovo ThinkShield security portfolio, extending existing EDR coverage into AI governance for ThinkShield endpoints. Enables security teams to identify sanctioned and shadow AI services in use across the workforce, and supports detection and response against prompt injection and AI misuse on Lenovo endpoints.

– Names the partnership and threats covered but no setup path [launch-20261021-4cb06267](#)

06 **Policy and execution governance for agents** NEW 50

Adds allow/block controls for which agents run, who uses them, and what they can reach, including approval workflows for high-risk tools, plus agentic investigation and governed response that automatically completes an investigation and executes a response within customer-defined boundaries.

– Describes governance scope but no concrete policy configuration shown

[launch-20261021-7d5f1e18](#)

07 **Singularity ASOC SOC suite launch** NEW 45

Introduces Singularity ASOC, a new integrated suite combining AI SIEM, Agentic Investigation, Hyperautomation, AI Data Pipelines, threat intelligence and IOC

management, and case management in a single offering on the Singularity Platform.

– Names six bundled components but no mechanism for how they interoperate

launch-20261021-2547d6a6

08 Agent Traces in Singularity Data Lake NEW 45

Adds Agent Traces: every agentic event is monitored and sent to Singularity Data Lake for visualization, retention, search, and use in dashboards or detection rules.

– Names the data destination but not how to query or alert on traces

launch-20261021-7d5f1e18

09 PartnerOne program expansion for AI security NEW 45

Opens a Singularity ADR early access program to certified PartnerOne partners, introduces an AI Recon Assessment Service for partners to lead AI security conversations with prospects, and expands PartnerOne with new Implementation Engineer certifications covering **SentinelOne Prompt Security**, endpoint and identity protection, **AI SIEM**, **AI Data Pipelines**, and **Hyperautomation**.

– Names the programs and certifications but not enrollment mechanics

launch-20261021-a0322cde

10 Singularity Credits consumption currency NEW 40

Introduces Singularity Credits, a unified consumption currency for AI-powered actions across the Singularity Platform, with administrator-configurable guardrails on AI investigation and automated response usage that bound the scope of every automated action to customer-defined limits.

– Explains the concept but not how limits are configured

launch-20261021-a0322cde

launch-20261021-2547d6a6

11 AI SIEM with built-in case management NEW 35

Adds AI SIEM with built-in incident and case management, surfacing threats from ingested security data so analysts start from prioritized findings rather than raw logs.

– States the benefit but not the underlying detection mechanism

launch-20261021-2547d6a6

12 Agentic Investigation of alerts NEW 35

Adds Agentic Investigation, which investigates alerts end-to-end using native and third-party security data, gathering evidence and working toward a verdict autonomously.

– Describes autonomous behavior but no scope limits or configuration

launch-20261021-2547d6a6

13 Wayfinder expert-led MDR services NEW 30

Extends Singularity ASOC with Wayfinder expert-led MDR, threat hunting, and incident response services for teams that want human analysts operating alongside the automated system.

– Names the service but not its scope or delivery model launch-20261021-2547d6a6

14 **Hyperautomation for response workflows** NEW

28

Adds Hyperautomation to close the loop from verdict to remediation, running routine response actions and workflows at machine speed.

– Generic description with no named workflow or trigger launch-20261021-2547d6a6

WAS THIS USEFUL?



Wiz Sensor

[Feedback](#)



[changelog](#) · Oct 20, 2026

[BLOG](#)

Wiz Sensor is a lightweight runtime agent from Wiz, the cloud security company, that collects workload telemetry from VMs, containers, and Kubernetes nodes to detect threats and vulnerabilities.

Wiz Sensor is now a native GKE add-on (Preview) enabled via `--`

`addons=WizSensor=ENABLED`, with fleet-wide Organization Policy enforcement and off-cluster secret management.

NEW DETECTION

NEW INTEGRATION

NEW SUPPORTED PLATFORM

Enables the Wiz Runtime Sensor as a first-party GKE add-on (Preview) via `--addons=WizSensor=ENABLED` on the `gcloud` CLI, covering both GKE Standard and GKE Autopilot clusters.

Supports fleet-wide enforcement through custom Google Cloud Organization Policies that automatically block unprotected cluster creation across projects, folders, or entire organizations.

Integrates with GKE's native Secret Manager and Workload Identity Federation so API keys and pull tokens are vaulted off-cluster and rotate automatically without cluster restarts.

Decouples sensor detection-rule and threat-definition updates from GKE cluster patch versions, node pool recreations, and workload maintenance windows.

WAS THIS USEFUL?



Iru

Feedback



changelog · Oct 22, 2026

CHANGELOG

RANK

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

Iru shipped Platform SSO OpenID sign-in for macOS 27, AI-generated BitLocker and Android Restrictions policies, and FileVault key auto-rotation, alongside a major expansion of Apple OS update controls covering rolling enforcement, phased rollouts, beta enforcement, split schedules, and hardware-aware compatibility checks.

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

01 OpenID sign-in for macOS 27 via Platform SSO NEW 91

The Single Sign-On Library Item adds OpenID sign-in support and web login settings for macOS 27 and later, enabling identity-provider web sign-in (including MFA) at the login window, Lock Screen, and FileVault unlock. Admins can sync the web login password to the local account and set a web login URL allow list.

Let users sign in to a macOS 27 device at the Lock Screen through your IdP's MFA web page instead of a local password.

📌 In the Iru web app, open the Single Sign-On Library Item › enable OpenID sign-in › set the web login URL allow list › enable 'Sync web login password to local account' › scope to macOS 27+ devices › save and deploy.

– Names exact settings and gives full deployment steps changelog-20261022-005d4688

02 FileVault key auto-rotation setting NEW 80

The FileVault Library Item gains an **Automatically rotate keys after viewing** setting that automatically rotates a device's FileVault recovery key after an administrator views it, with configurable rotation timing; key views remain recorded in the device activity timeline.

Ensure FileVault recovery keys are rotated immediately after an admin views one, closing the window where a viewed key could be reused.

📌 In the Iru web app, open the FileVault Library Item › enable **Automatically rotate keys after viewing** › set the rotation delay › save and deploy.

– Names config setting and exact UI deployment path changelog-20261022-005d4688

03 **Managed OS Library Items: new enforcement and rollout controls**

NEW

73

(new) Managed OS Library Items gain several new enforcement and rollout capabilities: hardware-aware OS enforcement that checks Apple's GDMF metadata service to skip incompatible updates on hardware-ineligible devices; a Rolling Enforcement option that delays sending a new target version declaration to Apple devices by up to 90 days; phased rollout scheduling with a configurable window of 1 to 7 days (set in hours) to spread OS declaration delivery; and beta OS enforcement using AppleSeed for IT beta tokens from Apple Business, with automatic token sync when Automated Device Enrollment (ADE) integrations are configured.

– Names the exact metadata service and enforcement modes Exact delay limit named, no config path given Exact window range named, mechanism described briefly Names token source and ADE sync mechanism

launch-20261020-01a75481

04 **AI-generated BitLocker and Android Restrictions policies** NEW

70

Iru AI now builds Library Items for BitLocker on Windows and Android Restrictions directly from a plain-language description, letting admins describe a policy (e.g. 'Require AES-256 encryption on all fixed drives and store recovery keys in Iru') and have Iru AI generate, review, and deploy the configuration.

Have Iru AI generate a BitLocker encryption policy Library Item from a plain-language description instead of manually configuring each setting.

📌 In the Iru web app, open Iru AI > describe the BitLocker policy you need (e.g. 'Require AES-256 encryption on all fixed drives and store recovery keys in Iru') > let Iru AI create the Library Item > review and deploy.

– Example prompt and deploy flow given, mechanism light changelog-20261022-005d4688

05 **Android Enterprise managed Google domain upgrade** IMPROVED

65

Admins can upgrade an Android Enterprise integration from a managed Google Play account to an organization's managed Google domain directly from the Iru web app, without re-enrolling devices.

– Names exact upgrade path and no-reenroll benefit changelog-20261022-005d4688

THINNER COVERAGE BELOW

06 **Separate target OS version in Automated Device Enrollment**

IMPROVED

55

Automated Device Enrollment gains a target OS version setting in addition to the existing minimum required OS version, allowing a separate enrollment floor and update target for macOS, iOS, and other supported platforms.

– Names the setting but no UI steps given `changelog-20261022-005d4688`

07 **Separate minor and major update schedules** IMPROVED

50

A single Managed OS Library Item can now set separate enforcement schedules for minor updates and major upgrades, eliminating the need for extra scoping logic.

– Describes the change but no concrete config detail `launch-20261020-01a75481`

08 **Recommended privacy permissions for Mac** IMPROVED

50

The Privacy Library Item's **Recommended privacy permissions** area now supports Mac (macOS apps and Safari website permissions), in addition to iPhone and iPad.

– Names config area and platforms, no steps `changelog-20261022-005d4688`

09 **Korean localization across Iru apps** NEW

35

Adds Korean localization to Passport, Self Service, Liftoff, and the Iru menu bar app in Iru Agent for Mac 5.2.7 (5401).

– Names apps and version but no further detail `changelog-20261022-005d4688`

WAS THIS USEFUL?  

syzkaller

Feedback



changes since 2026-10-21

CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller gains a linux-config-update agent skill and disables sandboxing for make configs.

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-21
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-21
```

Disables sandboxing when running `make configs` targets via the `Makefile`.

▼ 1 more release in this issue · 2026-10-21 → 2026-10-22

changes since 2026-10-20

CODE ↗

syzkaller extracts differential fuzzing logic into a reusable `pkg/fuzz` package with a public `ErrSkipFuzzing` sentinel

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-20
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-20
```

Exposes `fuzz.ErrSkipFuzzing` as a public sentinel in the new `syz-cluster/pkg/fuzz` package, allowing callers to detect and skip fuzzing sessions when base and patched kernel binaries are identical.

Extracts differential kernel fuzzing orchestration into a standalone `syz-cluster/pkg/fuzz` package (`fuzz.go`) reusable by `syz-cluster` workflows and local tools.

Extracts kernel build logic into `syz-cluster/pkg/build` for reuse across `syz-cluster` components.

Extracts smoke boot test logic into `syz-cluster/pkg/boot` for reuse across `syz-cluster` components.

Adds recognition of AUTOSEL tags and version-range stable backports in `syz-cluster/pkg/api`.

+1 more

WAS THIS USEFUL?



Edgescan

Feedback



Release Notes - 6th Oct 2026

NOTES

RANK

Edgescan is a managed vulnerability management and security testing platform that combines automated scanning with manual penetration testing to identify and help remediate security weaknesses in applications, networks, and infrastructure.

Edgescan improved sitemap management with a new unscoped-URL filter and manual sitemap uploads for pentested assets, and added inline scanner IP documentation on the asset page.

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

01 Manual sitemap upload for pentested assets NEW 60

Users can manually upload sitemap files (crawled and tested URL lists) for assets that are manually pentested, giving these pen-tests full sitemap coverage and visibility in the customer portal. The upload option is available directly on the asset page for the pentested asset.

When configuring a pen-tested asset, upload a crawled URL list so the customer portal reflects complete sitemap coverage without requiring an automated scan.

In the customer portal, navigate to the asset page for a manually pentested asset, then use the manual sitemap upload option to upload your crawled and tested URL list file.

– Names UI location and purpose but no file format limits Release Notes - 6th Oct 2026

THINNER COVERAGE BELOW

02 Unscoped URL filter on sitemap view NEW 47

A new filter on the sitemap view shows only URLs that have not yet been matched to a scope directive, making it easier to identify unscoped requests during target selection.

– Describes purpose but no exact UI path or command Release Notes - 6th Oct 2026

03 Scanner IP documentation on asset page IMPROVED 45

Scanner source IP documentation is now surfaced directly on the asset page, letting users allowlist scan traffic without needing to open the create-asset form.

– Clear benefit stated but no specific IPs or config detail Release Notes - 6th Oct 2026

WAS THIS USEFUL?



IDENTITY, ACCESS & DATA PROTECTION

Veeam Backup & Replication

Feedback



changelog · Oct 22, 2026

PRESS ↗

RANK

Veeam Backup & Replication is backup, recovery, and replication software for virtual, physical, cloud, and SaaS workloads.

Veeam expanded its iResOps intelligent recovery operations with a new Okta identity resilience Private Preview and opened a Private Preview of Veeam Data Platform built around a new DataAI Command Platform cloud console with AI agents, cross-estate search, and centralized access control; it also announced future previews for GitHub, Azure DevOps, and Google IAM.

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

01 DataAI Command Platform for Veeam Data Platform NEW

80

Veeam iResOps for Veeam Data Platform enters Private Preview, bringing contextual intelligence and precision recovery to on-premises and hybrid Veeam-protected workloads. It adds the **DataAI Command Platform** cloud console, which connects existing backup servers in minutes with no data migration or re-architecture required, the **DataAI Command Graph** mapping protected infrastructure, workloads, and indexed files in a single console, global search across protected and production data with full restore-point context, three built-in AI agents — Dashboard Agent, Graph Agent, and Backup Admin Agent — and centralized role-based access control (RBAC), single sign-on (SSO), and multi-factor authentication (MFA) management.

– Names console, graph, agents, and auth features precisely

2026-10-22-cae8e51b

2026-10-22-a93c4a73

02 iResOps identity resilience for Okta NEW

73

Veeam iResOps for Okta enters Private Preview, creating immutable restore points across 29 Okta object types — users, groups, applications, policies, and admin roles — each verified as restorable against the live Okta API. It enables granular comparison against a known-good state down to the individual property level, then restores objects and their relationships (memberships, app assignments, federation, roles) in the correct sequence, extending iResOps identity resilience already available for Active Directory and Microsoft Entra ID.

– Private Preview only; mechanism detailed but no hands-on steps

2026-10-22-a93c4a73

THINNER COVERAGE BELOW

03 Planned iResOps previews for GitHub, Azure DevOps, Google IAM

Veeam plans Private Previews in Q4 2026 for GitHub and Azure DevOps, covering code, pipelines, settings, and permissions, and for Google IAM, covering identity resilience across major cloud identity platforms.

– Only announced plans, no preview access yet 2026-10-22-a93c4a73

WAS THIS USEFUL?



Veeam Data Cloud

Feedback



changelog · Oct 22, 2026

PRESS >

Veeam Data Cloud is a SaaS backup and recovery service that protects Microsoft 365, Azure, and virtual machine data with cloud-hosted storage and management.

Veeam iResOps for M365 GA adds AI agent visibility via DataAI Command Graph, file-level undo, ROT data quarantine, coverage-gap detection, and plain-language query agents.

NEW INTEGRATION

Adds Veeam Intelligent ResOps (iResOps) for Microsoft 365 as a generally available add-on to Veeam Data Cloud for Microsoft 365, available in Foundation and Advanced editions and through Veeam Cloud & Service Providers (VCSPs).

Connects to the DataAI Command Graph, linking production Microsoft 365 data to backup context for unified visibility across data, access, AI agents, and backup protection.

Discovers known and unknown AI agents (including Microsoft Copilot) operating in the Microsoft 365 environment, maps which sensitive data each agent can reach, and identifies oversharing gaps.

Tracks AI-driven file changes file by file, enabling teams to identify what an agent modified or deleted and undo those changes precisely.

Identifies Redundant, Obsolete, and Trivial (ROT) data and enables quarantine of that data while teams decide what to archive or delete, reducing both blast radius and storage costs.

+2 more

WAS THIS USEFUL?



LoginRadius

Feedback



changelog · Oct 22, 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 strengthens token and passkey security with DPoP token binding, tighter passkey registration/verification controls, per-IP rate limiting across SSO and IAM APIs, and default Advanced Audit Logging V2 for new tenants.

←→ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 DPoP token binding for access and refresh tokens

NEW

60

Adds support for DPoP (RFC 9449) to bind Access Tokens and Refresh Tokens to the client's private key, preventing reuse of stolen tokens.

– Mechanism and standard named, but no config flag given `changelog-20261022-192c7d29`

02 Passkey registration and verification security controls

IMPROVED

60

Adds server-side checks to block duplicate Passkey registrations and enforce a per-account Passkey count limit. Also increases the Passkey verification challenge timeout from 1 minute to 3 minutes and enforces single-use, expiry-checked challenges to block replay attacks.

– Gives concrete before/after timeout change and behaviour `changelog-20261022-192c7d29`

THINNER COVERAGE BELOW

03 Per-IP rate limiting for SSO and IAM APIs

NEW

55

Adds per-IP rate limiting for SSO (OAuth 2.0, OIDC, and SAML), Account Management, and Partner IAM APIs.

– Names covered APIs but no limits or config exposed `changelog-20261022-192c7d29`

04 Advanced Audit Logging V2 enabled by default

IMPROVED

50

Enables Advanced Audit Logging (V2) and Team Audit Logs V2 by default for new Trial and Enterprise tenants.

– States which tenants get it, no further mechanism `changelog-20261022-192c7d29`

05 New default allowed callback domain NEW

35

Adds `lrcontent.com` to the default allowed callback domains.

– Single named domain, minimal context `changelog-20261022-192c7d29`

06 Faster JSON processing for core requests IMPROVED

30

Improves JSON processing speed for registration, profile update, and workflow requests.

– No metrics or mechanism given, generic performance claim `changelog-20261022-192c7d29`

WAS THIS USEFUL?



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 split its native Expo components into a standalone package to shrink app bundles, connected its MCP server directly to Clerk accounts for agent-driven setup, and tightened Secret key visibility in the Dashboard.

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

01 Native Expo components moved to separate package

BREAKING

91

Clerk's native Expo components (`<AuthView />` , `<UserButton />` , `<UserProfileView />`) are split into a dedicated `@clerk/expo-native-components` package, imported from `@clerk/expo/native-components` , so apps that don't use native components no longer bundle the iOS and Android SDKs; available in `@clerk/expo` 4.11.0. The `@clerk/expo/native` import path is deprecated and will be removed in the next major version — update imports after installing `@clerk/expo-native-components` . On Android, apps without the new package no longer migrate biometric credentials enrolled with earlier SDK versions (affected users must re-enroll), and the new app binary must ship before any over-the-air update containing this change or native components will throw a runtime error.

Migrate an Expo app that uses native Clerk components to the new separate package so the main bundle no longer includes the iOS/Android SDKs.

```
$ npx expo install @clerk/expo-native-components
```

– Names exact packages, import paths, version and migration/runtime risks

changelog-20261022-30e86c91

THINNER COVERAGE BELOW

02 Secret key masking after creation

IMPROVED

50

The `API keys` page in the Clerk Dashboard now reveals a Secret key's full value only at the moment of creation, showing only a partial value (first and last characters) thereafter to reduce the risk of key leakage.

:

– Clear UI change and location, but no config or API surface

changelog-20261022-30e86c91

03 Clerk MCP server account access

NEW

44

The Clerk MCP server now connects to your Clerk account, enabling agents to create applications, configure authentication, manage users, and set up production instances via your MCP client's sign-in flow.

– Describes capability but no commands or config shown changelog-20261022-30e86c91

L--> ALSO FROM THESE RELEASES

.

L--> BREAKING ON UPGRADE

- ! The `@clerk/expo/native` import path is deprecated and will be removed in the next major version; update imports to `@clerk/expo/native-components` after installing `@clerk/expo-native-components`.
- ! On Android, apps without `@clerk/expo-native-components` no longer migrate biometric credentials enrolled with earlier SDK versions — affected users must re-enroll. The new app binary must be shipped before any over-the-air update that includes this change, or native components will throw a runtime error.

WAS THIS USEFUL?



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

ConductorOne introduced a new Browse page that consolidates all requestable access—apps, profiles, AI tools, and actions—in one place, reorganizing navigation across the console.

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

01 Browse page consolidates requestable access

BREAKING

70

A new Browse page centralizes all requestable access into tabs for Apps, Profiles, AI tools, and Actions. App accounts that were previously on the Requests page move to **My resources > Accounts**, while the Requests page is refocused to task queues only — **Assigned to me**, **My requests**, and **Temporary** — opening by default on **Assigned to me**. Managers can still browse and request on behalf of another user via the existing **Manage** view, and bookmarks or links to the old Requests tabs redirect automatically to their new locations.

Request access to an AI tool for yourself or on behalf of a team member using the new Browse page.

In the ConductorOne portal, go to Browse > AI tools, find the tool you need, and click Request. To request on behalf of someone else, use the Manage for option.

Locate and request access to an AI tool without switching between tabs.

In the ConductorOne console, navigate to Browse > AI tools, find the tool you need, and submit an access request from that page.

– Names exact tabs, queues, and migration path; UI-only, no API surface.

Browse page (2026-10-21)

BREAKING ON UPGRADE

! Requestable access (previously on the Requests page) has moved to the new Browse page; existing bookmarks and links to the old Requests tabs are redirected automatically.

WAS THIS USEFUL?



Ping Identity

Feedback



20 Oct 2026 (2026-10-21)

[BLOG](#)

RANK

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

Ping Identity shipped a sandbox Remote MCP Server that lets AI agents administer Advanced Identity Cloud over OAuth 2.0, alongside PingOne updates for session cookie persistence control, team-scoped authorization modules, and JSON rendering in Forms.

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

01 Modules for grouping authorization resources in PingOne Authorize

NEW

78

Introduces modules in PingOne Authorize to group authorization resources by application team, letting teams self-manage their module's contents while administrators retain centralized governance. Authorization resources created before this feature are automatically migrated to a **Default** module and cannot be moved to other modules.

Organise authorization resources by application team in PingOne Authorize by creating a module so each team can self-manage their policies independently.

 In the PingOne console, go to Authorize > Modules, create a new module for your application team, and assign the relevant authorization resources to it.

Organize authorization resources by team in PingOne Authorize by creating a module and assigning relevant resources to it.

 In the PingOne console, go to Authorize > Modules, create a new module for your application team, then add the team's authorization resources to that module.

– Includes migration behavior and clear UI navigation steps [product docs](#)

02 Session token cookie persistence control

NEW

75

Adds per-environment control over whether PingOne uses a persistent or non-persistent session token (ST) cookie, configurable in Authentication Settings. Non-persistent cookies are discarded on browser close, useful for kiosk and shared workstation deployments.

Switch a shared kiosk environment to non-persistent ST cookies so sessions are cleared when the browser closes, preventing credential leakage between users.

 In the PingOne console, go to your environment > Settings > Authentication Settings, then set the session token cookie to non-persistent.

Configure non-persistent ST cookies for a kiosk environment so sessions are discarded when the browser closes.

 In the PingOne console, go to your Environment > Authentication Settings, then under **Session token cookie persistence** select **Non-persistent**.

– Gives UI path and mechanism but no API or flag [product docs](#)

03 Remote MCP Server for AI agent administration NEW

65

Launches an opt-in Remote MCP Server in sandbox, letting AI agents administer Advanced Identity Cloud via natural-language prompts over OAuth 2.0. Exposes tenant-scoped MCP tools covering journey and script configuration, application and identity provider management, user administration, policy settings, and audit log analysis.

– Lists tool scope but no setup command or endpoint [20 Oct 2026 \(2026-10-21\)](#)

THINNER COVERAGE BELOW

04 JSON value rendering in PingOne Forms IMPROVED

35

PingOne Forms now renders JSON-formatted values in dynamic text as plain text on end-user forms, enabling use of JSON outputs in flow development nodes.

– Describes behavior but no usage steps or config shown [product docs](#)

ALSO FROM THESE RELEASES

.
. .
.

BREAKING ON UPGRADE

! Authorization resources created before the modules feature are automatically migrated to the **Default** module and cannot be moved to other modules.

WAS THIS USEFUL?



1Password

[Feedback](#)[1Password for Android beta 8.12.44-24 \(2026-10-21\)](#)[BLOG](#)[RANK](#)

All the latest releases from the 1Password team

1Password's latest beta round renames its MCP server tools (a breaking change for automation), adds post-quantum passkey support on Android, and introduces a password-hiding toggle in the browser extension, alongside Watchtower 2FA naming clean-up and assorted iOS, Mac and Linux polish.

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

01 **Hide generated passwords toggle** NEW 75

Adds a **Hide generated passwords** toggle in Settings > Security to keep suggested passwords hidden in the inline menu, password generator, and extension popup.

Hide auto-suggested passwords from the inline menu and popup so they are never exposed on shared or observed screens.

In the 1Password extension, go to Settings > Security and enable **Hide generated passwords**.

– Exact setting path and effect, with a usage example.

[1Password in the browser 8.12.44-24 \(20...](#)

02 **MCP server tools renamed** BREAKING 70

The 1Password MCP server tools **list_variables** and **append_variables** are renamed to **list_secrets** and **append_secrets**; any automation, scripts, or integrations calling the old names will break.

– Names exact old and new tool names but no migration timeline.

[1Password for Linux beta 8.12.44-24 \(20...](#)

[1Password for Mac beta 8.12.44-24 \(2026...](#)

[1Password for Windows beta 8.12.44-24 \(...\)](#)

03 **Watchtower 2FA naming and QR scanning improvements** IMPROVED 70

Renames the Watchtower multi-factor authentication card to **2FA available** to match the passkeys card naming convention, relabels the **Scan QR code** option to **Scan 2FA QR code** for clearer intent, adds clearer wording and a direct link to 2FA setup instructions in the Watchtower **Use 1Password for 2FA** banner in item details, and extends the **Scan 2FA QR code** option so it works for any login, including from the Watchtower MFA banner.

– Names exact relabeled strings but no numeric scope.

1Password for Android beta 8.12.44-24 (...)

1Password for Linux beta 8.12.44-24 (20...

THINNER COVERAGE BELOW

04 **iOS UX and performance polish** IMPROVED 55

Improves performance when 1Password loads autofill suggestions, improves the search field when moving an item to another vault, shows **Enter a title** as a placeholder in the title field when creating a new item instead of the category name, and updates the SSO migration prompt link to read **Learn more about single sign-on**, opening end-user rather than administrator documentation.

– Four named tweaks but no metrics on the performance gain.

1Password for iOS beta 8.12.44-24 (2026...

05 **Post-quantum passkey support on Android** NEW 45

1Password for Android beta supports saving and signing in with post-quantum passkeys on sites that support them.

– States the capability but no setup steps or site list.

1Password for Android beta 8.12.44-24 (...)

06 **Autofill setup guidance for Touch ID prompt** IMPROVED 40

Adds a new Autofill setup step explaining how to disable the Touch ID prompt that appears on each fill.

– Clear setup-flow starting point but no exact UI path.

1Password for Mac beta 8.12.44-24 (2026...

07 **Automatic extension-app connection for Zen Browser** NEW 40

Adds automatic connection between the 1Password browser extension and the 1Password app when set up in Zen Browser.

– Names the browser but no mechanism for the connection.

1Password for Mac beta 8.12.44-24 (2026...

08 **Clear button and Escape shortcut in sign-in search** IMPROVED 35

Adds an X button and Escape key shortcut to clear the search field in the sign-in prompt.

– Names the UI control but is a minor, self-contained tweak.

1Password in the browser 8.12.44-24 (20...

09 **API Credential field relabeled to key name** IMPROVED

Relabels the first field in API Credential items from `username` to `key name`.

– Exact field names but purely cosmetic and unexplained.

1Password for Android beta 8.12.44-24 (...)

10 Improved security for passkeys shared across related sites

IMPROVED

30

Improves security of passkeys on sites that allow use of the same passkey across related websites.

– No mechanism or scope given for the security improvement.

1Password in the browser 8.12.44-24 (20...)

↳ BREAKING ON UPGRADE

- ! The 1Password MCP server tools `list_variables` and `append_variables` are renamed to `list_secrets` and `append_secrets` – any automation or scripts calling the old names will break.
- ! The 1Password MCP server tools `list_variables` and `append_variables` are renamed to `list_secrets` and `append_secrets` – any automation or integration calling the old names will break.

WAS THIS USEFUL?



RANK

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

Auth0 shipped Agent Mode for its CLI, a coordinated set of changes that make the CLI usable by AI agents: a dedicated flag for structured output, full command discovery, schema-validated payload submission, built-in docs search and AI-skill installation, structured JSON errors, and a keychain-free auth mode for CI.

←→ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 Schema and payload flags for commands NEW 86

`--schema` prints the exact payload a command expects, sourced live from the OpenAPI spec, and `--data` lets a caller submit an entire payload as a single JSON object that is validated locally before it reaches the API, removing guesswork around field names.

Let an agent discover the full CLI surface and create a resource safely in a single session, without guessing field names.

```
$ auth0 apps create --schema
auth0 apps create --agent-mode --data '{"name":"billing-
app","app_type":"spa"}'
```

— Both flags named with mechanism and a combined usage example

Agent Mode for the Auth0 CLI (2026-10-2...

02 Keychain-free auth mode for CI NEW 86

The `AUTH0_CLI_AUTH_MODE=env` environment variable, combined with `AUTH0_DOMAIN` and either `AUTH0_API_TOKEN` or client credentials, runs the CLI without touching the OS keychain or writing to disk, intended for CI runners and sandboxes.

Run the CLI in a keychain-free CI environment using only environment variables — no disk writes, no interactive prompts.


```
$ export AUTH0_CLI_AUTH_MODE=env
export AUTH0_DOMAIN=<your-auth0-domain>
export AUTH0_API_TOKEN=<management-api-token>
auth0 apps list --agent-mode --json-compact
```

– Exact env vars with a runnable CI example [launch-20261020-3342408b](#)

03 Agent Mode flag for AI-driven CLI use NEW

80

Adds a `--agent-mode` flag (or `AUTH0_AGENT_MODE=true` environment variable) that switches CLI output to structured JSON, disables interactive prompts and color, and guarantees non-zero exit codes on failure; it is auto-detected for recognized agent clients.



Filter a log list by query parameters without memorising per-command flags — useful for agents building dynamic queries.

```
$ auth0 logs list --agent-mode --query '{"per_page":5}'
```

– Names exact flag and env var and the behaviour change

[Agent Mode for the Auth0 CLI \(2026-10-2...](#) [launch-20261020-3342408b](#)

04 Docs search and AI skill install commands NEW

77

`auth0 docs search` queries Auth0's public documentation from the terminal without requiring login, returning raw-markdown URLs when run in agent mode. `auth0 agent skills install` installs the Auth0 skill into Claude Code, Cursor, or all supported AI assistants in one command, pre-loading Auth0 CLI conventions before first use.

– Both subcommands named with their exact behaviour

05 Structured JSON error responses IMPROVED

76

CLI errors now carry a stable `code` (e.g. `usage`, `auth`, `validation`, `not_found`, `rate_limit`), a finer `reason` sub-field (e.g. `local_validation`), a `details` array of per-field validation failures, and an HTTP `status` field for API-originated errors, replacing the previous free-text error sentence.

– Exact error field names and codes given, no direct command

06 Full CLI discovery via commands `--flat` NEW

75

`auth0 commands --flat` lists every runnable CLI command one per line for full discovery in a single call; appending `--detailed` returns complete invocation shapes, so an agent can build its action space without trial and error.

Discover every available CLI command and its full invocation shape in one call, useful for feeding an agent its action space.

```
$ auth0 commands --flat --detailed
```

– Exact subcommand and flags with a runnable example

↳ ALSO FROM THESE RELEASES

•
•
•

WAS THIS USEFUL?



Ory Network, Ory Hydra, Ory Kratos, Ory Keto, Ory Oathkeeper, Ory Polis v26.4.4 released

NOTES

RANK

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

Ory shipped external-ID matching for identity event streams and rolled out a redesigned Console navigation and layout across Ory Network.

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

01 External-ID matching for event stream events NEW 50

Event stream events — identity, login, registration, recovery, settings, verification, and session — can now be matched to users by external ID.

– Names event types but no field or API detail Ory Network, Ory Hydra, Ory Kratos, Ory...

02 Redesigned Ory Console layout and navigation IMPROVED 35

A new Ory Console layout and navigation rolled out across all Ory Network workspaces and projects, with no opt-in or migration required.

– Describes rollout scope but no specific UI paths or changes

The new Ory Console is available on Ory...

WAS THIS USEFUL?



Microsoft Entra

Feedback



changelog · Oct 19, 2026

BLOG ↗

RANK

Microsoft Entra is Microsoft's identity and access management suite, providing cloud-based sign-in, single sign-on, multifactor authentication, and access governance for users, apps, and devices.

Microsoft Entra expanded its Agent ID system to cover unmanaged and previously unknown agents, and introduced Dynamic Conditional Access that applies graduated, risk-driven enforcement to agent identities.

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

01 Dynamic Conditional Access for agents NEW 60

A new Dynamic Conditional Access capability enables risk-driven, graduated policy responses for agents: reducing permissions, blocking writes while preserving reads, narrowing data/resource scope, restricting network paths, shortening or revoking time-bound access, requiring human approval, or isolating/stopping agent runtimes, coordinated across multiple enforcement points.

– Explains graduated enforcement behaviors but no exact policy UI/config path

2026-10-19-a16a71ec

THINNER COVERAGE BELOW

02 Agent ID coverage for unmanaged and unknown agents NEW 45

Microsoft Entra Agent ID now extends to additional agent patterns, including local and previously unknown agents discovered through security sensors, making unmanaged agents into policy targets even without prior Microsoft Security integration.

– Describes scope and mechanism but no config surface or command 2026-10-19-a16a71ec

WAS THIS USEFUL?



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 shipped Agent Uninstall Protection with rotating maintenance tokens, a macOS 27 Golden Gate installer blocker, expanded Browser Discovery to third-party MDM-enrolled devices, an MCP Server toggle for AI client access to the User Portal, and new Android device policies.

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

01 macOS 27 Golden Gate installer blocker policy NEW

85

Adds the **Block macOS Golden Gate Installer** policy to prevent the macOS 27 Golden Gate installer from launching on MDM-enrolled Macs, automatically terminating the installer on launch.

Prevent macOS 27 Golden Gate from being installed on managed Macs before IT has validated the upgrade.

📌 In the Admin Portal, go to Devices > Device Policies > + Add New > Device Policy, select macOS, and choose the **Block macOS Golden Gate Installer** policy. Assign it to the relevant Mac device group.

— Exact policy name and effect, with UI assignment path `changelog-20261022-163d65ef`

02 Browser Discovery expansion via third-party MDM IMPROVED

85

Browser Discovery in AI & SaaS Management can now start when a user signs in to a JumpCloud-managed device, or when the JumpCloud Go extension is deployed via a third-party MDM provider using a **Client ID** and **Client secret** added to the JumpCloud Go extension distribution policy, expanding SaaS visibility without requiring full JumpCloud device management.

Deploy the JumpCloud Go extension through a third-party MDM to start browser-based SaaS discovery without full JumpCloud device management.

📌 In your third-party MDM console: 1. Go to AI & SaaS Management > Browser Discovery in the JumpCloud Admin Portal and copy the **Client ID** and **Client secret**. 2. Paste those values into the JumpCloud Go extension distribution policy in your MDM provider. 3. Deploy the policy. Discovery begins after the extension installs and the user signs in to the User Portal or a JumpCloud SSO application.

– Names exact credentials and policy location for deployment [product docs](#)

THINNER COVERAGE BELOW

03 **New Android device management policies** NEW 40

Adds new Android policies covering Display Settings, Autofill, Work Account authentication, and Bluetooth file sharing controls.

– Policy areas named but no config keys or steps provided [Release Notes 2026 \(2026-10-21\)](#)

04 **Customize Email and Automation Variables settings refresh** IMPROVED 25

Refreshes the **Customize Email** and **Automation Variables** tabs in Settings with shorter labels and improved helper text for easier navigation.

– Only a vague UI label/text tweak described [product docs](#)

↳ ALSO FROM THESE RELEASES

.
. .
.

WAS THIS USEFUL?



Cato Networks

Feedback



Make Your SIEM Smarter with Cato (2026-10-21)

BLOG ↗

RANK

Cato networks is a cato SASE Cloud converges SD-WAN, security, global backbone, and remote access into a global cloud-native service. WE ARE SASE™

Cato adds a Custom HTTP Push Connector with a data dictionary for flexible SIEM streaming, plus Cato XOps to correlate raw events into consolidated stories and surface east-west lateral movement alongside its growing library of native vendor integrations.

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

01 Cato XOps event correlation and east-west visibility

NEW

67

Cato XOps correlates raw events into consolidated, context-rich stories before forwarding to a SIEM, reducing ingestion volume — demonstrated compressing 37 million daily events into 94 stories. It also adds east-west WAN traffic visibility alongside north-south traffic, surfacing lateral movement signals across enterprise resources.

•
– Concrete compression metric but no config or command shown

Make Your SIEM Smarter with Cato (2026-...

02 Custom HTTP Push Connector for SIEM streaming

NEW

65

The **Custom HTTP Push Connector** streams selected Cato Events and Flows to any compatible HTTP destination, including AWS S3 and custom analytics platforms, with control over source, destination, filters, and context. A built-in data dictionary within the connector translates Cato telemetry field names into human-readable labels to ease SIEM field mapping.

•
– Names the connector and its config controls but no setup steps

Make Your SIEM Smarter with Cato (2026-...

THINNER COVERAGE BELOW

03 Expanded native vendor integrations

IMPROVED

32

Native integrations now extend to over 100 vendors, including CrowdStrike, Microsoft, and OpenAI, for enriching SIEM workflows with Cato SASE Platform telemetry.

– States a count and named vendors but no mechanism

Make Your SIEM Smarter with Cato (2026-...

WAS THIS USEFUL?



Sophos Protected Browser

Feedback



2026.44Released on October 21, 2026

NOTES ↗

RANK

Sophos Protected Browser is a secure, isolated browser from Sophos that runs web sessions in a sandboxed environment to protect endpoints from web-based threats and malicious content.

Sophos Protected Browser 2026.44 adds integration with Sophos DNS Protection, letting administrators apply browser-specific DNS threat policies.

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

01 Sophos DNS Protection integration for browser DNS

NEW

50

Sophos Protected Browser now integrates with Sophos DNS Protection to apply a dedicated DNS Protection policy to browser DNS requests, independently of device-level DNS settings, enabling separate threat controls for browser versus device activity.

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

2026.44Released on October 21, 2026

WAS THIS USEFUL?



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

Tailscale launched MagicRoute, a beta feature for Premium/Enterprise tailnets that automatically routes clients to the nearest healthy subnet router, app connector, or Tailscale Service host, alongside three new manual route-selection modes.

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

01 MagicRoute and expanded route-selection modes NEW

75

The admin console's `Settings > General > Route Selection > Configure` section gains `MagicRoute`, which gives each client automatic selection of the nearest healthy subnet router, app connector, or Tailscale Service host without relying on DERP regions, and supports tailnets with custom DERP servers (previously excluded from regional routing). The same Configure section adds three additional manual modes — `Active-passive failover`, `Regional routing`, and `Regional routing with in-region failover` — and clients on version 1.86.0 or later gain MagicRoute-based recommended exit node suggestions.

Enable MagicRoute for your tailnet so every client automatically connects to its nearest healthy subnet router or app connector — replacing the old regional routing toggle.

📌 In the admin console, go to `Settings > General > Route Selection > Configure` and select `MagicRoute`.

— Names config path, modes, and version but lacks failover mechanism details

launch-20261019-817d2a72

WAS THIS USEFUL?



Varonis

Feedback



announced · Oct 20, 2026

[BLOG](#) ↗

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 integrates with Collibra, enriching Collibra catalog assets with continuously updated sensitive data classifications for governance and AI use cases.

NEW INTEGRATION

Adds a Varonis integration with Collibra that sends sensitive data classifications — including PII, PHI, financial information, credentials, secrets, and regulated data — directly into the Collibra data catalog.

WAS THIS USEFUL?



GOVERNANCE, RISK & COMPLIANCE

Compyl

Feedback



Oct 20th, 2026

NOTES ↗

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

V2 Flat File Integration APIs add key/JWT auth and CRUD operations; Tasks gains saved views; policy same-owner/approver assignment now audited.

Adds V2 Flat File Integration APIs supporting `replace`, `append`, `parse`, `update`, and `delete` operations, with API key and JWT authentication; failed parses return an error without saving partial data.

The API Keys page now displays the environment's API base URL for quick endpoint discovery.

Adds saved views and filter persistence across the `Board`, `List`, and `Approvals` views in Tasks.

Policy administrators can now assign the same person as policy owner and approver, with a required warning, acknowledgment, and business justification recorded in the audit trail.

Policies attached to a control now count toward its pass/fail compliance evaluation.

+4 more

UPDATED INTEGRATION

Axonius integration expanded with additional asset inventory data

NEW INTEGRATION

New integration blueprints for recently onboarded data sources

⚠ BREAKING ON UPGRADE

! The `users` field is now required on IT Assets and Tasks API update calls; omitting it previously cleared role assignments silently but now returns a clear error.

WAS THIS USEFUL?



AGENT

Runlayer

[Feedback](#)[1.32.0 - Sharing API and MCP session stops \(2026-10-21\)](#)[BLOG](#)[RANK](#)

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

Runlayer shipped a beta Terraform provider for managing agents, schedules, connectors, and policies as code, a Sharing API for single-call principal access control, new MCP tools to stop and clear sessions (with a tightened validation rule), and infrastructure changes isolating session compaction on both hosted and self-hosted deployments.

←→ [WHAT SHIPPED](#) • [10 FEATURES](#) most completely described first [? what's the number?](#)

01 Terraform provider resources for agents, schedules, connectors, policies, and team management

[NEW INTEGRATION](#)[NEW](#)

98

Adds a beta Terraform provider (downloads.runlayer.com/runlayer/runlayer , version `~> 1.5`) for managing Runlayer as code: the `runlayer_agent` resource with `name` , `prompt` , `model` , `servers` (using per-connector `tool_names`), `egress_policy` (`default_action` and `allow` fields for per-agent network access), `env` , `tags` , `example_prompts` , and `is_disabled` ; the `runlayer_agent_schedule` resource (since provider `1.5.0`) with `agent_id` , `cron_expression` , `schedule_semantics` , and `timezone` ; the `runlayer_server` and `runlayer_policy` resources for connectors and policies; and, added in provider `1.8.1` , the `runlayer_team_member` and `runlayer_team_group_binding` resources for managing role and directory-group membership as code (applying either adopts an already-existing membership or binding). Adds `runlayer_server` and `runlayer_agent` data sources (the latter exposing `tool_count` , `skill_count` , `env_keys`) for CI checks, supports `terraform import runlayer_agent.<name> <uuid>` , and authenticates via `RUNLAYER_API_KEY` and `RUNLAYER_BASE_URL` environment variables. Full schema is documented under [Help → Terraform Reference](#) in the app.

Define an agent with least-privilege tool access and a network egress policy so security teams can review agent permissions like `IAM` in a pull request.

```

hcl

terraform {
  required_providers {
    runlayer = {
      source = "downloads.runlayer.com/runlayer/runlayer"
      version = "~> 1.5"
    }
  }
}

provider "runlayer" {}

resource "runlayer_agent" "release_notes" {
  name = "release-notes"
  prompt = trimspace(file("${path.module}/prompts/release-
notes.md"))
  model = "claude-sonnet-5"
  servers = [{
    server_id = runlayer_server.github.id
    tool_names = ["list_pull_requests", "get_pull_request"]
  }]
  egress_policy = {
    default_action = "deny"
    allow = ["github.com", "*.githubusercontent.com"]
  }
  env = {
    RELEASE_CHANNEL = "stable"
  }
}

```

Attach a recurring cron schedule to an agent so it runs automatically every Monday morning — change the day with a one-line diff and a reviewed PR.

```

hcl

resource "runlayer_agent_schedule" "sprint_digest" {
  agent_id = runlayer_agent.team["sprint-digest"].id
  name = "Sprint digest"
  input_prompt = "Summarize last week's closed issues for
the team."
  cron_expression = "0 9 * * 1"
  schedule_semantics = "local_wall_clock"
  timezone = "America/New_York"
}

```

Import an existing Runlayer agent into Terraform state by its UUID so all future changes go through plan and review.

```
$ terraform import runlayer_agent.support_triage 9f1c2b7e-...
```

Manage team membership and directory-group bindings as code in Terraform so role assignments are version-controlled and repeatable.

```
resource "runlayer_team_member" "analyst" {
  team_id = runlayer_team.soc.id
  user_id = "user_abc123"
  role    = "member"
}

resource "runlayer_team_group_binding" "soc_ad_group" {
  team_id = runlayer_team.soc.id
  group_id = "dir_group_xyz"
}
```

– Names every resource, field, and exact import/auth commands. Names resources, provider version, adoption behavior, and has an HCL example.

launch-20261019-3ee2c247

1.32.0 - Sharing API and MCP session st...

02 MCP session stop and clear tools

BREAKING

75

Adds `stop_session` and `clear_session_stop` tools to the Runlayer MCP so agents and automation (e.g. SOAR playbooks) can stop a session and clear a stop; requires the same permission as the API, a reason is mandatory, and the audit log records whether the action came from the API or the MCP. Stopping a session via the API now requires a session Runlayer has already recorded – an unknown session ID is refused rather than silently creating a stop record.

– Names both tools and the mechanism, but no runnable example given.

1.32.0 - Sharing API and MCP session st...

03 Session compaction isolated from ingest infrastructure

IMPROVED

65

Moves session storage compaction to its own ECS service so heavy compaction no longer interrupts new session activity in the Sessions view, and separately moves session compaction to its own Deployment on Kubernetes via the Runlayer operator for self-hosted installs, isolating heavy session work from ingest and allowing Kinesis installs to scale past one replica.

– Explains before/after architecture for two deployment targets, no user-facing action.

1.31.0 - Go deliverables and session la...

1.32.0 - Sharing API and MCP session st...

04 Layer-capture tags on session turns NEW

60

Adds layer-capture tags to each prompt, response, and tool card in Sessions, showing which layer (Agent runtime, LLM gateway, or Client hook) captured a turn when more than one layer recorded it.

Identify which layer captured each turn in a multi-layer session to debug duplicate turns during an Agent runtime + LLM gateway recording.

 In the Sessions view, open any recorded session → each prompt, response, and tool card now displays a layer tag (Agent runtime, LLM gateway, or Client hook) directly on the card.

– Names the three layers and gives a navigation path to see them.

1.31.0 - Go deliverables and session la...

THINNER COVERAGE BELOW

05 Go generates finished artifacts instead of describing them

IMPROVED

50

Go now builds and links the finished artifact — table, chart, one-page doc, or file — instead of describing it in chat when asked to compare options, show a trend, plan a trip, or send files.

– Describes before/after behavior with artifact types, no example.

1.31.0 - Go deliverables and session la...

06 Expanded PrivateLink connector capacity IMPROVED

40

Expands PrivateLink support for Runlayer Deploy connectors beyond four connections, with additional connections using AWS rule capacity shared across deployments.

– Explains capacity mechanism but no setup steps.

1.32.0 - Sharing API and MCP session st...

07 Bulk delete in Artifacts library NEW

40

Adds bulk delete to the Artifacts library, letting you select multiple artifacts and remove them in one step.

Clean up stale artifacts in bulk instead of deleting them one at a time from the library.

 In the Artifacts library, check the checkboxes next to the artifacts you want to remove → select **Delete selected** to remove them all in one step.

– Simple UI feature with steps but no further mechanism.

1.31.0 - Go deliverables and session la...

08 Sharing API for principal access control NEW

35

Adds a Sharing API that reads, sets, or removes one principal's access to an agent, skill, or plugin with a single idempotent call keyed by principal.

– No endpoint path, fields, or example given. [1.32.0 - Sharing API and MCP session st...](#)

09 Auto-review and Your review tabs in run drawer **NEW**

35

Adds Auto-review and Your review tabs to the run drawer, with draft notes preserved when switching between them.

– UI feature named but no navigation detail or example.

[1.32.0 - Sharing API and MCP session st...](#)

10 Build commit display in Settings **NEW**

30

Adds the exact build commit with a copy button to Settings → General next to the version number.

– Minor UI addition with exact location but no further detail.

[1.31.0 - Go deliverables and session la...](#)

NEW DETECTION [AI dictation app inventory in AI Watch Detect](#)

↳ **BREAKING ON UPGRADE**

- ! Stopping a session via the API now requires a session Runlayer has already recorded — an unknown session ID is refused rather than silently creating a stop record.

[1.30.0 - Eval replay and Go follow-through \(2026-10-21\)](#)

BLOG ↗

STABLE

Adds eval replay with recorded tool results, agent account OAuth connectors, improved Go follow-through, auto-saving run reviews, and an enhanced offline eval run picker.

Adds a **Replay recorded results** option under Tool mocking in agent evals, feeding each matching call its original recorded result so every replay sees identical data while remaining side-effect free.

Adds an OAuth connector workflow for agent accounts so admins can attach a connector without linking their own credentials — each user then grants access with their own account.

Go now continues completing unfinished parts of a request at the end of a turn instead of prompting whether it should, while still returning decisions involving money, access, data deletion, or messages to the user.

Run reviews in the run drawer now auto-save Pass/Fail, notes, and expected output after a pause in typing or on drawer close, replacing the explicit Save button with a **Saved** indicator.

Adds an offline eval run picker with search and filter by status, start date, or topic; shift-click range selection; quick picks for recent failures or topic spreads; and an Ask AI option that returns matching runs with reasons.

[+1 more](#)

WAS THIS USEFUL?



GOVERN

AI-Infra-Guard

Feedback



v4.6.7

NOTES ↗

RANK

A full-stack AI Red Teaming platform securing AI ecosystems via Agent Scan, Skills Scan, MCP scan, AI Infra scan and LLM jailbreak evaluation.

AI-Infra-Guard's v4.6.7 release expands Agent-Scan with CopilotKit protocol support, a cost-saving inference preset, and finer scan-input control, alongside new infrastructure fingerprinting for detection.

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

01 AG-UI (CopilotKit) protocol stream support in Agent-Scan NEW

45

Agent-Scan now supports AG-UI (CopilotKit) protocol streams, enabling scanning of CopilotKit-backed agents.

– Names protocol and backend but no usage steps v4.6.7

02 Prompt and skills forwarding to scanner IMPROVED

40

Agent-Scan now forwards the prompt and selected skills to the scanner, giving operators direct control over scan inputs.

– Explains benefit but no concrete interface named v4.6.7

03 Cheaper Inference provider preset for Agent-Scan NEW

35

Adds a Cheaper Inference provider preset to Agent-Scan for cost-sensitive scan configurations.

– Describes purpose but no config details v4.6.7

04 astron-agent fingerprint for infrastructure detection NEW

35

Adds **astron-agent** fingerprint to the data layer for AI infrastructure detection.

– Names fingerprint but no usage or scope detail v4.6.7

WAS THIS USEFUL?



ifixai

[Feedback](#)[v4.2.1](#)[NOTES ↗](#)

Independent Auditing of AI Agents. Run by human or the agent itself, to answer the most crucial question in the AI Agent Economy.

V11 JSON output uses `findings_only: true` and a `findings` array; report-only findings surface at the top of all output formats; `ifixai compare` judges V11 by verdict; all judge votes are retained in multi-judge runs.

↳ [GET THIS VERSION](#)

```
$ git clone --branch v4.2.1 https://github.com/ifixai-ai/iFixAi.git
# already have the repo? check out this version:
$ git checkout v4.2.1
```

NEW DETECTION V11 JSON output carries `findings_only: true` and a `findings` array in place of `score`, `threshold`, `passing`, and confidence interval fields, reflecting its report-only status.

Report-only findings now appear at the top of markdown, HTML, JSON, and terminal output.

`ifixai compare` now compares V11 results by verdict rather than by percentage.

Every judge's vote is now retained in multi-judge runs, including judges that returned no verdict.

UPDATED DETECTION [Inspection catalogue expanded to 61 inspections](#)

WAS THIS USEFUL?



Docker Sandboxes

[Feedback](#)

announced · Oct 21, 2026

[BLOG](#) ↗

RANK

Docker Sandboxes run AI coding agents in isolated environments on your machine or on Docker-managed cloud infrastructure

Docker Sandboxes added configuration controls for bypassing the proxy on cloud SSH connections, blocking in-sandbox browser launches and OAuth redirects, and tuning daemon log rotation.

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

01 Daemon log rotation controls NEW

78

Three new config keys govern daemon log rotation: `daemon.log.maxSizeMB` sets the MB threshold at which the daemon log rotates (minimum enforced), `daemon.log.maxBackups` caps the number of rotated logs kept (`0` removes the count limit while the age limit still applies), and `daemon.log.maxAgeDays` sets retention in days (`0` removes the age limit while the backup count limit still applies).

Tune daemon log rotation so logs rotate at 50 MB, keep at most 5 backups, and are purged after 30 days — whichever limit is hit first.

```
yaml
daemon.log.maxSizeMB: 50
daemon.log.maxBackups: 5
daemon.log.maxAgeDays: 30
```

– All three config keys and their interactions named [product docs](#)

02 Bypass upstream proxy for cloud SSH NEW

75

A new `proxy.sshDirect` config setting lets cloud SSH connections skip the upstream proxy entirely while other host CLI requests continue routing through the proxy; set it to `false` to re-enable the proxy for SSH. Useful when a network proxy blocks the SSH gateway port or tunnels.

When your network proxy blocks the SSH gateway port, bypass it for cloud SSH while keeping other CLI traffic routed through the proxy.

```
yaml
proxy.sshDirect: true
```

Route cloud SSH connections directly, bypassing the upstream proxy, without restarting the daemon — useful when the proxy blocks SSH tunnels.

```
yaml  
  
proxy.sshDirect: true
```

– Named config key with exact behaviour and override value

[product docs](#)

03 Block sandbox browser launches and OAuth redirects

NEW

65

A new `browser.enabled` config key prevents sandboxed processes from opening URLs in the host browser and blocks in-sandbox OAuth logins that redirect to localhost.

Lock down a sandbox so it cannot open host browser tabs or trigger OAuth redirect flows during automated runs.

```
yaml  
  
browser.enabled: false
```

– Named config key, behaviour clear, no deeper mechanism given

[product docs](#)

NEW INTEGRATION BAND AI kit for Docker Sandboxes

announced · Oct 21, 2026

[BLOG](#)

Docker Sandboxes adds `sbx secret set-custom` and `sbx policy allow network` flags plus a `band.yaml` kit config to connect isolated sandbox agents to multi-agent BAND rooms.

Adds `docker.io/bandhq/band-python-kit:latest` (current published tag `6.0.0`) as a Docker Sandbox kit, enabling sandboxed agents to connect to BAND AI rooms via outbound WebSocket with one `sbx create` command.

Supports `band.yaml` config file in the sandbox workspace to declare which agent entrypoint (`agent.id`, Python file) the BAND kit launches on `sbx run`.

Adds `sbx secret set-custom` with `--sandbox`, `--host`, `--env`, and `--placeholder` flags to scope per-sandbox credentials (e.g. `BAND_API_KEY`, `ANTHROPIC_PROXY_KEY`) to specific outbound hosts without exposing real keys inside the microVM.

Adds `sbx policy allow network --sandbox <name> <host>` to grant a named sandbox outbound access to specific API hosts (e.g. `api.anthropic.com`).

Enables multi-agent BAND room composition where sandboxed local agents and hosted agents collaborate via @mentions, shared message history, and tool-

execution events visible to human participants.

WAS THIS USEFUL?



CrowdStrike Falcon AI Detection and Response

Feedback



changelog · Oct 21, 2026

PRESS ↗

CrowdStrike Falcon AI Detection and Response is a cybersecurity platform that uses artificial intelligence to detect, investigate, and automatically stop threats across endpoints, cloud, and identity environments.

Falcon Guardian AIDR extends AI agent security from software to silicon via NVIDIA OpenShell and planned NVIDIA Sentry on BlueField-4 DPU integration.

NEW INTEGRATION

Integrates Falcon Guardian AIDR with NVIDIA OpenShell secure runtime software to correlate agent prompts, model interactions, and runtime activity with broader Falcon platform security context for detecting malicious behavior across AI agent workflows.

NEW INTEGRATION

Plans in-silicon enforcement by integrating Falcon Guardian AIDR with NVIDIA Sentry on BlueField-4 DPUs for out-of-band monitoring and hardware-level policy enforcement to detect and contain unsafe agent behavior beyond the software runtime.

WAS THIS USEFUL?



PLATFORMS & UTILITIES

GitHub CLI

[Feedback](#)

v2.103.0

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

GitHub's official command line tool

GitHub CLI's biggest change this window is opt-in expiring, auto-refreshing OAuth credentials across the auth command suite, alongside a new (private preview) `gh issue artifact` command group for managing issue documents from the terminal.

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

01 Expiring, auto-refreshing authentication credentials

NEW

95

`gh auth login --expiring` (requires CLI v2.103.0+) requests short-lived access tokens paired with rotating refresh tokens to reduce leaked-token risk; `gh config set expiring_credentials enabled` opts every future sign-in into expiring credentials without passing the flag each time. `gh auth refresh --expiring` preserves expiring credentials when changing scopes (otherwise a scope change silently replaces an expiring credential with a non-expiring one). `gh auth token --no-refresh` prints the current token without triggering a refresh, and `gh auth status` gains `--json` (combinable with `--show-token`) plus `--no-refresh`, and now also displays whether a credential is refreshable and the relative time until the current access token expires.

Enable expiring, auto-refreshing OAuth tokens globally so every future `gh auth login` uses short-lived credentials without manual re-authentication.

```
$ gh config set expiring_credentials enabled
```

Opt in to expiring credentials for a single sign-in and configure git to use `gh` as the credential helper so expiring tokens are served and refreshed for git operations.

```
$ gh auth login --expiring && gh auth setup-git
```

Sign in with expiring credentials to limit the blast radius of a leaked token — the access token expires within hours and refreshes automatically.

```
$ gh auth login --expiring
```

Opt in to expiring credentials globally so every future `gh auth login` requests them without needing the flag.

```
$ gh config set expiring_credentials enabled
```

– Flags, config key and config file-equivalent all named with runnable examples v2.103.0

02 `gh issue artifact` command group for managing issue documents

NEW

90

New `gh issue artifact` command group (private preview) with `list`, `view`, `create`, `edit`, `download`, and `delete` subcommands for managing issue artifacts from the CLI. `gh issue artifact create` gains `--attach` to attach images or video files to Markdown documents, and `gh issue artifact edit` gains `--body-file` to update content from a file and `--restore-version` to roll back an artifact to an earlier version.

Create a Markdown artifact on issue #142 with an attached image, then list all artifacts on that issue to confirm it was created.

```
$ gh issue artifact create 142 signin-plan.md --attach ./signin-flow.png
gh issue artifact list 142
```

– Full subcommand and flag list with runnable example v2.103.0

03 Per-host targeting flag for auth commands IMPROVED

70

Adds `--hostname` to `gh auth setup-git`, `gh auth token`, and `gh auth logout` to target a single host rather than all authenticated hosts — for configuring the Git credential helper, reading a stored token, or signing out without interactive prompts respectively.

– Names three commands and the flag's effect on each product docs

04 Environment variables for Enterprise and default host authentication

NEW

65

Adds `GH_ENTERPRISE_TOKEN` for authenticating to GitHub Enterprise Server in non-interactive and GitHub Actions contexts, and `GH_HOST` to set a default host for commands not tied to a specific repository.

– Named env vars but no runnable example given product docs

05 Git credential helper via gh auth setup-git IMPROVED

60

Supports configuring GitHub CLI as a Git credential helper via `gh auth setup-git`, enabling `git push` and `git pull` without a separate credential manager.

– Named command and effect, no further mechanism given [product docs](#)

THINNER COVERAGE BELOW

06 gh extension browse rebuilt as TUI IMPROVED

30

Replaces `gh extension browse` with a Bubble Tea terminal UI.

– Bare mention, no behaviour or navigation detail given [v2.103.0](#)

WAS THIS USEFUL?



// END OF TRANSMISSION ■

The Cyber Toolchain · curated daily · October 22, 2026

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

► [Subscribe](#)

[Feedback on this issue](#)