THE CYBER TOOLCHAIN

NO. 150

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 25, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 280 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 71 tools

▾ // THE BRIEF HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 Opus 5

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

Three unrelated things stand out today: Aikido Security can now autofix through private registries and safer base images, SailPoint extends

vim ? help

governance to machine and agent identities, and Netcap cuts noisy network detections by counting per actor.

PROTECT

Fix dependency and image risk without leaving private supply chains

Aikido Security's AutoFix can now work with private package registries and `.aikido` repo context, while its zero-CVE base images give teams a safer starting point. SonarQube's Remediation campaigns let the agent batch backlog issues into one pull request, reducing the review overhead of fixing many code findings one by one.

Aikido Security · SonarQube

GOVERN

Put owners, approvals, and reviews around non-human access

SailPoint adds owner assignment, credential visibility, bulk actions, and approval workflows for non-human identities, including AI agents and machine accounts. ConductorOne and JumpCloud add more targeted access certification, so reviewers can audit profiles, entitlements, rules, enrollments, and compliance campaigns instead of chasing access evidence manually.

SailPoint · ConductorOne · JumpCloud

DETECT

Make network detections less noisy and easier to route

Netcap reworked all 37 built-in rules to count per actor and separated alerts from inventory, which helps analysts see behavior instead of repeated asset facts. rsigma adds per-tenant rulesets, STIX/TAXII-driven detections, enrichers, and faster batched delivery, making shared detection infrastructure easier to run without mixing tenant logic.

Netcap

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Autonomous Offensive Frameworks (AOF)

xalgorix

Feedback ⓘ v4.6.208 NOTES ↗ RANK

> xalgorix shipped two configuration options this window: a compact checklist mode that halves system-prompt token usage, and a smart stub mode that preserves security-relevant context when tool results age out.

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

01 Compact checklist mode for reduced prompt tokens

Not checked yet

NEW

93

Setting `XALGORIX_CHECKLIST=compact` (or choosing `compact` in Settings → Scan checklist) cuts per-call system-prompt tokens roughly in half, from ~27K to ~14K, while retaining the full 22-phase methodology, every vulnerability class, and every verifier; only inline command scripts are trimmed.

Persist the compact checklist across all scans by adding it to the environment file xalgorix reads on startup.

ini

```
XALGORIX_CHECKLIST=compact
```

— Exact env var, UI path, and before/after token numbers given v4.6.206

02 Smart bounded-context stubs for archived tool results

Not checked yet NEW

85

Enabling `XALGORIX_TOOL_ARCHIVE_SMART_STUB` (opt-in, off by default) keeps reflected payloads, error strings, secrets and tokens, HTTP status and notable headers, and endpoints in the agent's working context when tool results age out, while the full original result stays archived and retrievable.

Persist the setting across runs by adding it to the xalgorix env file so every scan session inherits it without re-exporting.

ini

```
XALGORIX_TOOL_ARCHIVE_SMART_STUB=true
```

— Names exact env var, default state, and what's retained v4.6.208

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

NetSPI OCInferno

Feedback ⓘ changes since 2026-11-19 CODE ↗

OCInferno adds cross-tenancy ENDORSE graph edges, DRG route table and Network Firewall policy object enumeration, and richer NAT Gateway and DRG Attachment fields.

Model cross-tenancy ENDORSE statements in the OpenGraph as `HAS_ENDORSE` edges from `OCIGroup` nodes to `OCIExternalTenancy` nodes, surfacing the source-tenancy side of cross-tenancy access pairs. Not checked yet

NAT Gateway enumeration now also collects `nat_ip` and `public_ip_id` fields, with `nat_ip` promoted to a default display column. Not checked yet

NEW INTEGRATION DRG Attachment enumeration now also collects the `drg_route_table_id` field. Not checked yet

NEW INTEGRATION Network Firewall policy objects enumeration: NetworkFirewallPolicyObjectsResource, network_firewall_address_lists, network_firewall_services, network_firewall_service_lists

WAS THIS USEFUL? 👍 👎

Rapid7 Metasploit Pro

Feedback ⓘ

November 24, 2026: 5.1.1-2026112301 NOTES ↗ RANK

Metasploit Pro is a penetration testing platform that automates vulnerability exploitation, network scanning, and security assessments to help professionals identify and remediate weaknesses.

> Metasploit Pro 5.1.1-2026112301 adds six new exploit and auxiliary modules across five product targets and modernizes an existing post-exploitation module to use a newer internal API.

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

01 **New modules for Plex, Zimbra, changedetection.io, Welcart, Atlassian DC**

Not checked yet NEW 65

Six new exploit and auxiliary modules were added targeting changedetection.io, Plex, WordPress Welcart, Zimbra, and Atlassian Data Center. A new auxiliary module performs arbitrary file read against Plex Media Server versions up to 1.43.2.10687 and the first 1.43.3 build 10828. A new exploit module targets Zimbra's administrator-authorized DeployZimletRequest operation when server-side Zimlet JSP execution is enabled. Details for the changedetection.io, WordPress Welcart, and Atlassian Data Center modules were not specified beyond the product names.

— Mechanism given for 2 of 6 modules; others named only by product

November 24, 2026: 5.1.1-2026112301

02 **dns_srv_lookup post module modernized** Not checked yet IMPROVED

65

`modules/post/multi/gather/dns_srv_lookup` was updated to use the new modern `create_process` API.

— Exact module path and API named, no behavioral detail given

November 24, 2026: 5.1.1-2026112301

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

Hak5 WiFi Pineapple Pager

Feedback ⓘ 2.1.0-stable

NOTES ↗ RANK

WiFi Pineapple Pager is a pocket sized WiFi Pineapple built for those on the go.

> WiFi Pineapple Pager 2.1.0 adds a kick-client API and command, buzzer audio playback via `pwmplay`, a USB-C storage recovery page, and improved handshake capture and client tracking, alongside SSH, configuration validation and network reconfiguration hardening.

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

01 **USB-C storage recovery page** Not checked yet NEW 90

Adds a storage recovery page over USB-C at `http://172.16.52.1:1471` to delete files or factory reset the device when the eMMC fills up; settings and installed software now live on the internal eMMC, with new firmware unpacked there before first boot.

Recover storage space or factory reset the Pager when the eMMC is full and the device is unreachable over Wi-Fi.

> 📍 Connect the Pager via USB-C, then open `http://172.16.52.1:1471` in a browser to delete files or trigger a factory reset from the recovery page.

— Exact URL, mechanism, and step-by-step recovery instructions. 2.1.0-stable

02 **Kick-client API and command** Not checked yet NEW 80

Disconnect clients from the Pager's APs, optionally blocking them from rejoining, using the `PINEAPPLE_KICK_CLIENT` command or the new kick API endpoint.

Disconnect a misbehaving client from the Pager's APs and block it from rejoining during an engagement.

```
$ PINEAPPLE_KICK_CLIENT <client_mac> block
```

— Names command and gives runnable example, but endpoint path not specified. 2.1.0-stable

03 **Buzzer audio playback via pwmplay** Not checked yet NEW 75

Play mp3 files through the Pager's buzzer as a standard ALSA device using the new `pwmplay` command.

Play an audio file through the buzzer to test alert tones or voice prompts during a field deployment.

```
$ pwmplay /path/to/alert.mp3
```

— Concrete command with runnable example. 2.1.0-stable

04 **Recon handshake capture and client tracking improvements**

Not checked yet IMPROVED 70

Evil WPA AP captures now raise the handshake alert for PMKID captures, and every client joining the Pager's WPA APs leaves a complete handshake, with management APs excluded from captures. A new Captures panel on the Virtual Pager's Recon page collects handshakes in one place, and clients of the Pager's own APs are now tracked by the specific SSID joined and the time they left, with every connection raising the client connected alert.

— Describes mechanism and UI location but no exact config or command. 2.1.0-stable

THINNER COVERAGE BELOW

05 **Settings validation and VPN config hardening** Not checked yet

IMPROVED 50

Wi-Fi passwords, SSIDs, country codes, hostnames, time zones, DNS hosts, and VPN settings are now validated before saving, and OpenVPN config is restricted to root-only read access.

— Lists validated fields but no UI path or error details. 2.1.0-stable

06 **Targeted network reconfiguration** Not checked yet IMPROVED

50

Network changes now reconfigure only what changed, leaving the management link and other radios up while a single AP, client mode, or VPN setting is updated.

— Explains behaviour change but no specific command or flag. 2.1.0-stable

07 **SHA-512 password hash support** Not checked yet NEW 40

SHA-512 password hashes are now supported, with stricter hash checks and constant-time comparison.

— Names the hash algorithm but gives no configuration steps. 2.1.0-stable

08 **SSH now uses dropbear** Not checked yet IMPROVED 35

SSH now uses dropbear, preserving existing host keys and `authorized_keys` across upgrades.

— States the change but no migration steps or config detail. 2.1.0-stable

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

Invicti

Feedback ⓘ changelog · Nov 25, 2026 CHANGELOG ↗ RANK

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

> Invicti rolled out agentic pentesting upgrades this window: internal agent tunneling for unreachable targets, two new attack agents (secrets detection and business logic), custom assessment profiles, authentication reporting, and richer Octo finding data and vulnerability-chain testing.

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

01 **Allow incomplete agentic assessments via env var** Not checked yet

NEW 75

Start assessments that allow an incomplete run when credits fall short by setting `INVICTI_ALLOW_INCOMPLETE_ASSESSMENTS` in Scan CLI environments; Octo warns that the assessment may be incomplete.

— Names exact env var and behavior but no full workflow shown changelog-20261125-e098b147

02 **Custom assessment profiles and rerun settings in Pentests**

Not checked yet NEW 75

On the Pentests settings page, create custom assessment profiles by setting `maximum duration`, `maximum credits`, and selecting which agents to include, instead of choosing only predefined profiles. Also introduces the Rerun with discovered credentials setting, which lets Octo restart testing with authentication when valid credentials are found during an attack (off by default).

— Names exact settings and UI location for configuration changelog-20261125-e098b147

03 **Automations triggers for agentic pentest completion**

Not checked yet NEW 70

Automate responses to agentic pentests by triggering on the new Agentic assessment completed and Agentic assessment failed events in Automations, with email notification support.

Trigger a Slack or email alert in Automations whenever an agentic pentest finishes or fails, so the security team is notified without polling the UI.

📍 In the Invicti Platform, go to Automations › New Automation, set the trigger to `Agentic assessment completed` or `Agentic assessment failed`, then add a `Send email` action and save.

— Names exact trigger events and includes a runnable setup example

changelog-20261125-e098b147

04 **Internal agent tunneling for unreachable targets** Not checked yet

NEW

70

Assess targets not reachable from the internet by assigning an internal agent to a target in Inventory; Octo tunnels all assessment traffic, including attack agent requests, through that agent.

Scan an internal web application that has no public internet exposure by routing all pentest traffic through an already-deployed internal agent.

📍 In Invicti Platform, go to Inventory › select your internal target › Edit › assign an Internal Agent, then start a new pentest from the New pentest wizard.

— Clear mechanism plus exact navigation steps to use it changelog-20261125-e098b147

05 **New Attack Agents** Not checked yet NEW

65

Adds two new attack agents: a secrets detection agent that scans client-side JavaScript and source maps (including original source code) for exposed credentials, uses AI to suppress publishable keys and surface actionable secrets, and when source code is provided also scans that source code and reports only credentials Octo can confirm grant access to the application; and a business logic agent that runs after all other attack agents complete and looks for business logic vulnerabilities they did not report.

— Detailed mechanism but no direct usage path given States behavior but no concrete scope or example

changelog-20261125-e098b147

THINNER COVERAGE BELOW

06 **Authentication report for login sequence assessments**

Not checked yet NEW

50

Download an authentication report for assessments that use a login sequence recording, showing each replayed login step, its screenshot, and matching HTTP activity for troubleshooting.

— Describes contents but no exact download path given changelog-20261125-e098b147

07 **Vulnerability chain agent out-of-band and source code support**

Not checked yet IMPROVED

50

The vulnerability chain agent now supports out-of-band testing and can use provided source code, marking findings inconclusive when the issue cannot be reproduced against the live target.

— Explains new testing modes and fallback behavior clearly changelog-20261125-e098b147

---

08 **Richer Octo finding metadata** Not checked yet IMPROVED

30

Octo findings now include a description, an impact statement, and a classification, also surfaced in Invicti Platform Inventory.

— Brief statement of added fields with no further detail changelog-20261125-e098b147

WAS THIS USEFUL? 👍 👎

◆ Code Security

cve-lite-cli

Feedback ⓘ v1.43.0 NOTES ↗ RANK

> cve-lite-cli v1.43.0 adds experimental Python requirements scanning alongside its JavaScript scanning, extends JSON output with ecosystem metadata, adds a file/stdout export for override audits, and ships several breaking changes to baseline paths, exit codes and Yarn fix-command counting.

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

01 Python requirements scanning alongside JavaScript

Not checked yet

NEW

85

Pinned Python requirements files can now be scanned alongside JavaScript lockfiles in a single run via `--python`; flags `--check-overrides`, `--check-maintenance`, `--check-licenses`, `--usage`, `--only-used`, `--offline`, and `--offline-db` are accepted but apply only to the JavaScript part. Discovery under `--python` finds every pinned requirements file in the project — any `*requirements*.txt`, `constraints*.txt`, and `.txt` files under a `requirements/` folder at any depth — and follows `-r` and `-c` includes the way pip reads them.

Scan a project that has both JavaScript lockfiles and Python requirements files in one pass, with override and license checks applied to the JavaScript part.

```
$ cve-lite /path/to/project --python --check-overrides --check-licenses
```

— Names all flags and discovery rules with a runnable example v1.43.0

02 Override audit export via --output

Not checked yet NEW

75

`cve-lite overrides --json --output <dir|file|->` writes override audit results to a file or directory; `--output -` sends clean JSON to standard output, for piping into tools such as a SIEM.

Export override audit results as clean JSON piped to another tool, such as a SIEM or reporting script.

```
$ cve-lite overrides --json --output -
```

— Runnable command with exact flag syntax given v1.43.0

03 **Exit code changes for incomplete scans and prod-only mode**

Not checked yet

BREAKING

70

Under `--incomplete-policy error`, combining `--python` with a JavaScript project where a Python gap exists (unpinned line or unreadable `-r` / `-c` include) now exits 3 instead of 0. A single-folder `--prod-only` scan whose lockfile contains only dev dependencies now exits 0 with a message instead of exiting 3.

— Explains before/after exit codes with named flags, no example v1.43.0

04 **Python baseline file path change** Not checked yet BREAKING

65

`--python` ratchet baselines are now written to `.cve-lite/baseline-pypi.json`; any existing Python baseline must be renamed from `baseline.json` to that path or it will be ignored.

— Exact path given with clear migration step v1.43.0

05 **Ecosystem metadata in JSON output** Not checked yet IMPROVED

60

`--json` output now carries an `ecosystem` field (`npm` or `pypi`) on every finding; multi-folder JSON also gains `units` and `notes` fields when Python requirements are scanned.

— Names fields but no example of the output v1.43.0

THINNER COVERAGE BELOW

06 **Yarn workspace fix-command coverage change** Not checked yet

35

BREAKING

In Yarn workspaces, fix commands change and the count of findings covered by a direct fix can decrease, because a finding also reached transitively is no longer counted as covered by the direct fix alone.

— Describes behavior change but names no flag or command v1.43.0

└──▷ **BREAKING ON UPGRADE**

! `--python` ratchet baselines are now written to `.cve-lite/baseline-pypi.json`; rename any existing Python baseline from `baseline.json` to that path or it will be ignored.

! Under `--incomplete-policy error`, combining `--python` with a JavaScript project where a Python gap exists (unpinned line or unreadable `-r` / `-c` include) now exits 3 instead of 0.

! A single-folder `--prod-only` scan whose lockfile contains only dev dependencies now exits 0 with a message instead of exiting 3.

! In Yarn workspaces, fix commands change and the count of findings covered by a direct fix can decrease, because a finding also reached transitively is no longer counted as covered by the direct fix alone.

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Nov 24, 2026 BLOG ↗ RANK

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

> Snyk's Evo platform gained a Global Risk Database, a searchable catalog of every AI model Snyk has scored—including models not yet in your repositories—so teams can compare risk and capability before adopting a model.

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

01 **Global Risk Database in Evo** Not checked yet NEW 79

Evo now includes the Global Risk Database, a catalog of every AI model Snyk has scored, including models not present in your repositories, so you can evaluate risk before adopting a model. It offers a risk-against-capability view plotting each model's overall risk score against its Epoch Capabilities Index, side-by-side comparison of two to five models at once by risk category and attacker goal, search and filters to find models by name, provider, or cost tier, model pages showing a model's full risk profile and capabilities, and Evo chat now answers questions about models in the catalog.

Before adopting a new LLM in your codebase, compare it side-by-side with current models by risk category and attacker goal to make an informed decision.

> 📍 In Evo, go to Global Risk Database › Side-by-side comparison, select two to five models, and review results by risk category and attacker goal.

Quickly surface candidate models by provider or cost tier before reviewing their full risk profile.

> 📍 In Evo, go to Global Risk Database, use Search and filters to filter by provider or cost tier, then open a model page to view its full risk profile and capabilities.

— Names concrete views and UI navigation but no API or config surface 2026-11-24-8b73792a

WAS THIS USEFUL? 👍 👎

OpenQodex

Feedback ⓘ openqodex@0.24.0 NOTES ↗ RANK

OpenQodex is a join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> OpenQodex 0.24.0 moves AI reviewer configuration from personal dotfiles to a shared, ordered project config with automatic fallback between reviewers, and updates the CLI override flags and GitHub Action to match.

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

01 Team-shared ordered reviewer list with fallback

Not checked yet

93

BREAKING

Configure an ordered list of reviewers in `review.reviewers` inside `.openqodex/config.yaml`, each entry specifying a tool (`claude` or `codex`) with an optional model and effort, shared across the whole team; when one entry hits a usage or rate limit, the review restarts with the next entry in the list. The `init` command now writes a `review.reviewers` list with `min_version: 0.24.0` into the project config when none exists. The personal keys `reviewer`, `reviewer_model` and `reviewer_effort` in `~/.openqodex/config.yaml` are no longer read, so teams must move this configuration into the project's `.openqodex/config.yaml`; a review made under a different reviewer list no longer satisfies the push hooks.

Define a team-shared reviewer priority list so that if Claude hits a rate limit, the review automatically retries with Codex.

yaml

```yaml
review:
  reviewers:
    - tool: claude
      model: claude-opus-5-5
      effort: high
    - tool: codex
      model: gpt-6.1-sol
```

— Names config key, file, fields and a runnable config example openqodex@0.24.0

02 CLI flags to override reviewer per run

Not checked yet

BREAKING

80

Pick a single reviewer for one run with `--reviewer <tool>:<model>` and `--reviewer-effort`, which replace the old `--reviewer-model` flag.

Override the reviewer for a single run — useful for quick checks or cost control — without changing the shared project config.

```
$ openqodex review --reviewer codex:gpt-6.1-sol --reviewer-effort low
```

— Exact flag syntax with a runnable command example `openqodex@0.24.0`

THINNER COVERAGE BELOW

03 GitHub Action reviewer config resolution

Not checked yet · IMPROVED

50

The GitHub Action now reads the reviewer config from the pull request's base commit instead of the base branch tip, and can run Codex as the reviewer when the runner is logged in.

— Describes behaviour change but no config or command to act on `openqodex@0.24.0`

04 Project-level web reviewer toggle

Not checked yet · BREAKING

40

`review.reviewer_web` in `.openqodex/config.yaml` replaces the old personal `reviewer_web` key, moving control of the AI reviewer's web-tool access from the user's personal config to the project level.

— Names the config key but gives no mechanism or example `openqodex@0.24.0`

└──▷ **BREAKING ON UPGRADE**

! The personal keys `reviewer`, `reviewer_model` and `reviewer_effort` in `~/.openqodex/config.yaml` are no longer read; move reviewer configuration into the project's `.openqodex/config.yaml` under `review.reviewers`.

! The `--reviewer-model` flag is gone; use `--reviewer <tool>:<model>` with `--reviewer-effort` instead.

! A review made under a different reviewer list no longer satisfies the push hooks.

WAS THIS USEFUL? 👍 👎

syzkaller

Feedback ⓘ changes since 2026-11-23 CODE ↗

Syzkaller is an unsupervised, coverage-guided kernel fuzzer developed by Google that automatically finds bugs in operating system kernels, primarily Linux.

syzkaller adds INVEPT and INVVPID support to x86 SYZOS nested virtualization and tightens AI job concurrency limits

Adds `SYZOS_API_NESTED_INTEL_INVEPT` and `SYZOS_API_NESTED_INTEL_INVVPID` API calls to x86 SYZOS to enable fuzzing of Intel nested virtualization TLB invalidation instructions. Not checked yet

Adds `bubblewrap` and `socat` to syzkaller Docker containers, making those tools available inside the fuzzing environment. Not checked yet

Displays the git revision on the `syz-hub` status page to make it easier to identify which build is running. Not checked yet

Limits concurrent auto repro-C and AI cause-bisection jobs via a shared `maxAutoJobsInProgress` map, replacing separate per-workflow rate-window logic. Not checked yet

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

More predictable notifications on GitHub.com (2026-11-24) BLOG ↗ RANK

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

> GitHub overhauled how notification reasons are determined and surfaced, introducing a fixed priority order for displayed reasons, multi-reason email headers, real-time team-membership awareness for team mentions, and a narrower Subscriptions page scoped to manual subscriptions only.

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

01 Fixed-priority notification reason ordering Not checked yet

70

BREAKING

GitHub now determines the notification reason shown in the inbox and email footer using a fixed priority order — `mention` > `team mention` > `assign` > `review requested` > `comment` > `manual` > `author` > `state change` > `subscribed` — rather than the most recently occurring activity type. This is a breaking change: email filters or automations keyed on a specific reason value may now match differently for existing threads.

— Full priority list given but no API/config surface to act on

More predictable notifications on GitHu…

02 Subscriptions page scoped to manual subscriptions Not checked yet

65

BREAKING

The Subscriptions page now shows only threads you explicitly subscribed to manually; use the notification inbox to manage and bulk-unsubscribe from threads where you are author, commenter, assignee, or reviewer. Workflows that relied on the Subscriptions page to bulk-manage those threads must move to the notification inbox.

— Names the exact UI pages and migration path More predictable notifications on GitHu…

03 All-reasons notification email headers Not checked yet BREAKING

60

Notification email headers now include all applicable notification reasons (such as `mention` and `subscribed`) rather than a single reason, so email filters and automations can act on every reason that applies to a thread. This is a breaking change for filters or automations that expected only one reason header value.

— Describes header mechanism but not exact header name

More predictable notifications on GitHu…

THINNER COVERAGE BELOW

04 **Real-time team membership for team-mention notifications**

Not checked yet IMPROVED 50

Team-mention notifications now reflect current team membership in real time: leaving a team stops future notifications from that mention, and joining a team enables notifications for future activity on already-mentioned threads.

— Clear before/after behaviour but no config surface

More predictable notifications on GitHu…

05 **No resubscribe after manual unsubscribe** Not checked yet IMPROVED 45

Routine activity such as commenting, submitting a pull request review, changing a thread's state, or merge queue activity no longer resubscribes you after a manual unsubscribe.

— States the fix but no mechanism detail beyond trigger list

More predictable notifications on GitHu…

06 **Team-mention reason for security advisory comments**

Not checked yet IMPROVED 32

Team mentions in security advisory comments now surface with a `team mention` reason instead of `commented`.

— Single-line behaviour change with no further context

More predictable notifications on GitHu…

└──▷ **BREAKING ON UPGRADE**

! Notification reason shown in inbox and email footer may change for existing threads: the displayed reason now follows a fixed priority order (mention > team mention > assign > review requested > comment > manual > author > state change > subscribed) rather than the most recent activity. Email filters or automations keyed on a specific reason value may match differently.

! Notification email headers now include ALL applicable notification reasons rather than a single reason. Filters or automations that expected only one reason header value may now match on additional reasons.

! The Subscriptions page no longer lists threads where you receive notifications due to authoring, commenting, assignment, mentions, or review requests — only manually subscribed threads appear. Workflows that relied on the Subscriptions page to bulk-manage those threads must move to the notification inbox.

WAS THIS USEFUL? 👍 👎

Codacy

Feedback ⓘ announced · Nov 24, 2026 BLOG ↗ RANK

Documentation homepage for the Codacy automated code review tool.

> Codacy launched a new Remediation report that gives security teams a dedicated view into SLA compliance, mean time to remediate, and backlog aging across all finding sources.

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

01 Remediation report for SLA tracking

Not checked yet NEW

88

A new Remediation report page under Security and risk → Reports → Remediation shows, per severity (Critical, High, Medium, Low): Within SLA or Over SLA status, mean time to remediate (MTTR) vs. the SLA, share of findings closed on time with counts, a weekly trend chart, and an open backlog split into 1–30, 30–90, and 90+ days overdue. It can be filtered by Segments, Repositories, and Scan types, with filter state encoded in the page URL so a shared link reproduces the exact view. SLAs are configured directly from the report via the Configure SLAs control at top right, defaulting to 30, 60, 90, and 120 days from Critical to Low when none are set. Coverage spans Codacy's own analysis, app scanning, container scanning, penetration tests, and security tickets imported from Jira; the report covers a rolling 90-day window, and longer audit records (e.g. six or twelve months) require capturing monthly snapshots via the Codacy API.

⋮

Navigate to the Remediation report to get a one-screen answer on SLA compliance before a quarterly security review.

> 📍 In Codacy, go to Security and risk → Reports → Remediation. Select the desired Segments, Repositories, and Scan types — the resulting URL encodes your filters so you can share the exact view with leadership.

— Names exact nav path, config control, defaults, filters, and data sources.

launch-20261124-ec758b11

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ announced · Nov 24, 2026

BLOG ↗ RANK

SonarQube is an enforce consistent quality gates across every pull request and pipeline, from IDE to CI/CD, with automated code review for bugs and security issues.

> SonarQube Cloud's biggest addition this window is Remediation campaigns, which let the Remediation Agent work through batches of backlog issues and submit one consolidated pull request, alongside bring-your-own-key LLM support for the agent.

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

01 **Remediation campaigns for backlog issues** Not checked yet NEW

88

SonarQube Cloud adds Remediation campaigns, a mode where the Remediation Agent works through up to 200 related backlog issues in steps on a separate campaign branch, then opens one final pull request for team review. Enable it via Manual backlog remediation under `Administration > AI capabilities > Remediation Agent` for GitHub-bound projects on Team (annual) or Enterprise plans. Scope a campaign by selecting specific issues, using Assigned to me, targeting issues Blocking the quality gate, or leaving all checkboxes cleared to let the agent start with the most severe eligible issues; provide a plain-language goal of up to 1,800 characters under Specify instructions for the agent. Track progress per step — including skipped, closed, and merged steps — from the project's Agent activity page filtered by the Campaign source. Campaigns support Java, JavaScript/TypeScript, Python, C, C++, C#, PHP, and secrets, and can include eligible SonarQube Hunter Agent findings regardless of language.

Enable Remediation campaigns for GitHub-bound projects so the Remediation Agent can work through backlog issues automatically.

📍 In SonarQube Cloud, go to Administration › AI capabilities › Remediation Agent, then under Enable agent toggle on Manual backlog remediation, and select Save.

Start a focused campaign to clear quality-gate-blocking issues in a project, with a plain-language goal to guide the agent.

📍 On your project's Issues page (main branch), select Blocking the quality gate under Recommended issues to be fixed, enter a goal such as 'Fix null pointer issues in the billing module' (up to 1,800 characters) under Specify instructions for the agent, then select Create campaign.

— Full mechanism, scope, limits and exact UI path given. launch-20261124-bd20379b

02 **BYOK LLM API keys for Remediation Agent** Not checked yet NEW

70

Add a bring-your-own-key (BYOK) LLM API key under `Administration > AI capabilities > LLM API Keys`, supporting OpenAI or Anthropic keys, up to three keys per organization; SonarQube Cloud verifies the key with the provider before saving.

Add a BYOK OpenAI or Anthropic API key so the Remediation Agent uses your own LLM provider instead of the Sonar-provided one.

> 📍 In SonarQube Cloud, go to Administration › AI capabilities › LLM API Keys, choose OpenAI or Anthropic, enter your key and a name, and select `Add key`. SonarQube Cloud verifies the key with your provider before saving.

— Named config path, providers and key limit, but no further mechanism.

launch-20261124-bd20379b

WAS THIS USEFUL? 👍 👎

Black Duck Feedback ⓘ

Polaris release update: Secure every language, automate dynamic testing, and meet develo

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 release adds LLM-powered Signal AI Code Analysis for any programming language, native work-item export to external issue trackers, configurable token expiration, and scheduling for fAST Dynamic scans, alongside Bridge 4.6.0 updates for scoped automated remediation and Rapid Scan Static.

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

01 Issue-type filters for Fix PRs and PR Comments Not checked yet

NEW 90

Restrict Bridge 4.6.0 automated remediation to specific finding types using the `polaris.fixpr.filter.issueTypes` and `polaris.prComment.filter.issueTypes` config keys, letting teams scope automated Fix PRs (e.g. to SCA) and PR Comments (e.g. to SAST) independently so noisy suggestions from one issue type don't trigger unwanted actions in the other.

Restrict automated Fix PRs to SCA findings only in a Bridge 4.6.0 config, so noisy SAST suggestions don't generate unwanted pull requests.

```yaml
polaris.fixpr.filter.issueTypes: SCA
```

Restrict PR Comments to SAST findings only while keeping SCA fix PRs separate, giving different teams independent control over automated remediation scope.

```yaml
polaris.prComment.filter.issueTypes: SAST
```

— Named config keys with runnable example values Polaris release update: Secure every la…

02 Signal AI Code Analysis for any language Not checked yet NEW

83

New Signal AI Code Analysis uses LLM reasoning to find vulnerabilities in any programming language — including COBOL, Bash, and RPG — without build configuration, surfacing findings as an AI tool type alongside existing SAST and SCA results. AI

Analysis can be enabled at the Organization, Application, or Project level to control rollout scope, and a custom LLM can be configured for Signal in Bridge 4.6.0.

Initiate an AI scan of a COBOL or Bash codebase that existing SAST rules don't cover, using the Source Upload workflow in the Polaris UI with no build configuration required.

> 📍 In the Polaris UI, go to your project, select Source Upload, upload your source archive, then view findings under the `AI tool type` filter alongside your SAST and SCA results.

— Mechanism and scope levels named; exact enable flag missing

Polaris release update: Secure every la…

---

03 **Work item export to external issue trackers** Not checked yet NEW

75

Export Polaris findings as native work items (Bugs, Tasks, or Issues) to Azure DevOps Server (2019, 2020, 2022), GitHub Enterprise Server (3.16–3.18), or GitLab Self-Managed (16.11–18.4) through a secure tunnel with bidirectional status sync.

— Names three trackers and exact supported versions, no setup command

Polaris release update: Secure every la…

---

04 **Configurable expiration for tokens** Not checked yet NEW

65

Set token expiration for service account tokens and personal access tokens (PATs) with preset durations (seven days, 30 days, 90 days), a custom date, or no expiration when creating or regenerating a token.

— Lists concrete presets, no exact UI path given Polaris release update: Secure every la…

---

05 **Scheduled fAST Dynamic scans** Not checked yet NEW 60

Schedule fAST Dynamic scans on a daily or weekly Test Scheduling Policy alongside SAST and SCA, with support for policy inheritance and auto-assignment during project onboarding.

— Describes scheduling and inheritance, no config key named

Polaris release update: Secure every la…

THINNER COVERAGE BELOW

---

06 **Rapid Scan Static (Sigma) workflows in Bridge** Not checked yet

NEW 25

Enable Rapid Scan Static (Sigma) workflows in Bridge 4.6.0.

— Bare mention with no mechanism or config detail

Polaris release update: Secure every la…

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Veracode CLI (2026-11-24) BLOG ↗ RANK

Veracode is a cloud-based application security platform that scans code for vulnerabilities using static, dynamic, and software composition analysis testing methods.

> Veracode's CLI gained five new subcommands (static fix automation, DAST login recording, Git-integrated static scanning, cached findings access, and toolchain discovery), broad packaging support for more languages and CI/CD signaling, RBAC for container/IaC scans, and an asynchronous remediation mode for SCA.

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

01 veracode package command enhancements

Not checked yet IMPROVED

87

The `veracode package` command's `veracode package discover` subcommand was extended to detect toolchains for Java projects using Gradle and JavaScript projects using npm or Yarn, and `veracode package` itself gained the `--csharp-source` flag to package C# source directly for Static Analysis, `--skip-packagers` to exclude artifacts such as unsupported SQL, `--java-source` and `VERACODE_PACKAGE_JAVA_SOURCE=true` for Java source packaging, `VERACODE_PACKAGE_GOLANG_GENERATE` for Go CGO code generation, `SRCLR_MAKE_WINDOWS` to preprocess C/C++ source with Make using a native Windows gcc or clang compiler, new exit codes `5`, `6`, and `7` for CI/CD pipeline automation, support for the Poetry package manager in Python autopackaging, and support for `requirements-dev.txt` in Python packaging workflows.

Package a C# project without a build step so you can submit source directly for Static Analysis in a CI pipeline that lacks a full .NET build environment.

```
$ veracode package --csharp-source
```

— Names command and detected toolchains, no example Lists every flag and env var verbatim with one runnable example

Veracode CLI (2026-11-24)

02 veracode dynamic record for DAST login scripts

Not checked yet

NEW

78

New `veracode dynamic record` command records a Chrome login sequence and saves it as a SIDE file that can be used as a DAST login script, useful when manual auth flows are hard to replicate.

Record a DAST login sequence in Chrome to generate a SIDE file you can upload as a login script, useful when manual auth flows are hard to replicate.

```
$ veracode dynamic record
```

– Explains mechanism and output format with runnable command Veracode CLI (2026-11-24)

03 **veracode fix static command with remote remediation**

Not checked yet NEW 73

New `veracode fix static` command requests Static Analysis remediation, with a `--remote` flag to trigger GitHub remediation requests directly, e.g. for use in a CI workflow without a local environment.

Trigger Static Analysis remediation from a GitHub Actions workflow using the remote flag so fixes are requested without a local environment.

```
$ veracode fix static --remote
```

– Names flag and exact command but limited mechanism detail Veracode CLI (2026-11-24)

04 **RBAC and findings handling in veracode scan** Not checked yet

IMPROVED 65

The `veracode scan` command added role-based access control (RBAC) for container and IaC scans, restricting developers to application profiles associated with their teams; categorizes secret-type findings as Critical; prioritizes vulnerability results from the NVD database when available; and allows applying a custom policy defined in the Veracode Platform, with policy evaluation for individual findings.

– Describes several behaviours but no flags or commands shown Veracode CLI (2026-11-24)

THINNER COVERAGE BELOW

05 **Fire-and-Forget async mode for veracode fix sca** Not checked yet

IMPROVED 55

The `veracode fix sca` command now supports asynchronous Fire-and-Forget (FnF) mode, submitting remediation requests for background processing without waiting for fixes or pull requests.

– Names the mode and effect but no flag or example Veracode CLI (2026-11-24)

06 **veracode scan static for Git-integrated scanning** Not checked yet

NEW 50

New `veracode scan static` command runs static scans and reviews findings within a Git workflow.

— Names command but no flags or example Veracode CLI (2026-11-24)

07 **veracode findings static for cached results** Not checked yet NEW

48

New `veracode findings static` command accesses cached static scan findings stored in the local Veracode directory.

— Names command and behaviour but no example given Veracode CLI (2026-11-24)

08 **Target UUID support in veracode dynamic scan** Not checked yet

IMPROVED

40

The `veracode dynamic scan` command now supports target UUIDs for triggering scans in addition to webhook UUIDs.

— Brief statement with no example or parameter detail Veracode CLI (2026-11-24)

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

aguara

Feedback ⓘ v0.29.0 NOTES ↗ RANK

> Aguara's v0.29.0 release tightens its capability-based findings — cutting false positives on explicit privilege drops, downgrading Kubernetes SETUID/SETGID requests to LOW severity, and restoring a regressed cron-write detector — while raising the Go toolchain requirement and expanding the analyzer/category vocabulary.

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

01 Kubernetes SETUID/SETGID severity downgrade

Not checked yet

95

BREAKING

Kubernetes workloads requesting `SETUID` or `SETGID` now report as a LOW observation under `K8S_IDENTITY_CAPABILITIES_001` instead of a HIGH finding, so they no longer fail a `--fail-on high` gate on their own; use `--fail-on low` to keep gating on these. Previously accepted `SUPPLY_007` findings for these requests reappear as `K8S_IDENTITY_CAPABILITIES_001` once the baseline is regenerated.

Re-gate CI on Kubernetes SETUID/SETGID requests after the severity downgrade, so a workload requesting those capabilities still fails the pipeline.

```
$ aguara scan --ci --fail-on low .
```

Regenerate your project baseline after upgrading so stale `MCP_009` and `SUPPLY_007` entries are replaced with the new finding IDs.

```
$ aguara audit . --write-baseline
```

— Exact finding ID, severity change, and re-gating command given v0.29.0

02 Context-aware MCP and supply-chain capability checks

Not checked yet IMPROVED

85

Capability drops, `setuid: false`, `setgid: false`, descriptions, and comments no longer trigger `MCP_009` or `SUPPLY_007`, reducing false alarms when a manifest explicitly drops privileges. Baselines must be regenerated with `aguara audit . --write-baseline` since stale `MCP_009` and `SUPPLY_007` entries accepted under the old logic no longer match.

Regenerate your project baseline after upgrading so stale `MCP_009` and `SUPPLY_007` entries are replaced with the new finding IDs.

```
$ aguara audit . --write-baseline
```

— Names exact flags, finding IDs, and the regeneration command v0.29.0

THINNER COVERAGE BELOW

03 New analyzer and category values for JSON findings

Not checked yet

BREAKING

55

JSON finding consumers that filter by `analyzer` or `category` should now also accept `mcp-capability`, `kube-policy`, and `container-security`.

— Lists exact new values, no integration example v0.29.0

04 Cron-write detection restored under SC-EX-007

Not checked yet

IMPROVED

52

Direct Python cron writes to `/etc/cron.d/<name>` or `/etc/crontab` are now detected again under `SC-EX-007`.

— Names paths and finding ID but no usage example v0.29.0

05 Go 1.26.0 toolchain requirement

Not checked yet

BREAKING

40

Building from source, `go install`, and importing the Go library now require Go 1.26.0 or later.

— Simple version bump, minimal elaboration v0.29.0

└──▷ BREAKING ON UPGRADE

! Building from source, `go install`, and importing the Go library now require Go 1.26.0 or later.

! A `SETUID` / `SETGID` capability addition no longer fails `scan --ci` or `--fail-on high` on its own; use `--fail-on low` to keep failing on these requests.

! Baselines must be regenerated with `--write-baseline`: `MCP_009` findings accepted in an older baseline no longer match, and accepted `SUPPLY_007` findings for Kubernetes capability requests reappear as `K8S_IDENTITY_CAPABILITIES_001`.

! JSON finding consumers that filter by `analyzer` or `category` must now also accept `mcp-capability`, `kube-policy`, and `container-security`.

WAS THIS USEFUL?  

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.167.15 NOTES ↗ RANK

Infisical is an open-source secret management platform that helps teams securely store, sync, and manage API keys, environment variables, and other sensitive credentials.

> Infisical's v0.167.15 release adds GCP workload identity federation for app connections, recursive folder access queries, and tightens API security with new response headers and admin sign-up restrictions.

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

01 **CSP and Permissions-Policy headers on API responses**

Not checked yet IMPROVED 60

Infisical's API now sends Content-Security-Policy and Permissions-Policy headers on all responses.

– Names exact headers, verifiable via response inspection v0.167.15

THINNER COVERAGE BELOW

02 **GCP workload identity federation for app connections**

Not checked yet NEW 50

Infisical adds GCP workload identity federation support as an authentication method for app connections.

– Names the integration but no setup steps provided v0.167.15

03 **Token-gated super-admin sign-up** Not checked yet NEW 45

Super-admin sign-up can now be restricted to users who hold an activation token, via a new token-gated sign-up flow.

– Describes the restriction but no exact flag or config key given v0.167.15

04 **Recursive folder access queries and redesigned access sheet**

Not checked yet IMPROVED 45

Recursive folder access queries from the root are now supported, paired with a redesigned folder access sheet that shows role context and recursion.

– Explains behaviour and UI change but no exact navigation path v0.167.15

05 **Re-import deleted certificates in PKI module** Not checked yet

NEW 40

The PKI module now allows re-importing previously deleted certificates.

— Clear capability named but no steps or limits given v0.167.15

06 **Reusable hooks on agent vaults** Not checked yet NEW 25

Reusable hooks can now be attached to agent vaults, one per service.

— Bare mention with no mechanism or configuration detail v0.167.15

WAS THIS USEFUL? 👍 👎

GitGuardian Detection Engine

Feedback ⓘ

Detection Engine Updates Version 2.180 (2026-11-20) BLOG ↗

GitGuardian Detection Engine scans code repositories and developer workflows to automatically detect exposed secrets, API keys, and credentials in real time.

NEW DETECTION CrowdSec LAPI Secret detector: CrowdSec LAPI secrets

NEW DETECTION Deno Organization Token detector: Deno Deploy organization access tokens

NEW DETECTION GAP Agent Token detector: GAP Cloud agent bearer tokens

NEW DETECTION instana_api_token detector: IBM Instana API tokens

NEW DETECTION idira_adaptive_authentication detector: Idira (formerly CyberArk) Adaptive Authentication

NEW DETECTION Lovable API Key detector: Lovable API keys

NEW DETECTION Qdrant API Key detector: Qdrant API keys

NEW DETECTION LINE Messaging API channel credentials checker

UPDATED DETECTION Alchemy API Key detector prefix update: Alchemy API keys

UPDATED DETECTION HashiCorp Consul ACL Token detector instance URL capture

UPDATED DETECTION humeai_api_key checker Hume Platform support

UPDATED DETECTION PostHog Personal API Key detector format update: PostHog personal API keys

NEW DETECTION Analyzers for 26 secret types: Cloudant Credentials, Cloudant IAM, Confluent Keys, GitGuardian Internal Monitoring Key, GitHub Enterprise Token and 21 more

UPDATED DETECTION Cloudant IAM analyzer account ownership reporting

UPDATED DETECTION IBM Cloud key analyzer identity and roles reporting

NEW DETECTION Auth0 revoker: Auth0 Keys

NEW DETECTION Supabase service_role JWT revoker: Supabase service_role JWTs

UPDATED DETECTION aws_iam revoker error reporting: AWS IAM Keys

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

DataDome

Feedback ⓘ Flutter Dio SDK 3.6.0 (2026-11-24) BLOG ↗

RANK

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

> DataDome's Flutter Dio SDK gained Android Gradle Plugin 9 support, and the Traefik module now forwards more header data for bot-detection analysis.

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

01 Android Gradle Plugin 9 support in Flutter Dio SDK

Not checked yet IMPROVED 65

The Flutter Dio SDK adds support for Android Gradle Plugin (AGP) 9, including its built-in Kotlin support, requiring Flutter 3.47+ and Kotlin 2.2.20+, while retaining AGP 8 support via the `kotlin-android` Kotlin Gradle plugin.

— Names exact versions and plugin but no usage steps Flutter Dio SDK 3.6.0 (2026-11-24)

THINNER COVERAGE BELOW

02 Larger header list truncation in Traefik module Not checked yet

IMPROVED 50

DataDome's Traefik module now truncates the `HeadersList` at 768 characters instead of 512, allowing more header data to be forwarded for bot-detection analysis.

— States exact config field and before/after limits Traefik module 1.4.0 (2026-11-24)

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Unfurl

Feedback ⓘ v2026.11.24 NOTES ↗ CODE ↗ RANK

Unfurl is an open-source forensic tool that deconstructs URLs, breaking them into components to reveal embedded metadata, timestamps, tracking parameters, and other hidden information.

> Unfurl's latest release overhauls its JSON output into proper nested and flat formats with regex filtering, adds new ID parsers for Timeflake, Notion, and Kimi IDs, and introduces safety limits for short-link expansion.

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

01 Real JSON output formats for CLI

Not checked yet IMPROVED

95

The CLI now emits real JSON via `-t json` (a nested tree where each node holds its derived nodes) or the new `-t json-flat` (a flat list using a `parent_id` field), both written as JSON Lines to stdout or via `-o`. Each output object includes `input`, `unfurl_version`, and a `lineage` field; previously the output was Python repr rather than valid JSON.

Extract all timestamps from a URL as machine-readable JSON, piped into jq for filtering — now that `-t json` outputs valid JSON Lines instead of Python repr.

```
$ unfurl -t json "https://example.com/?t=1700000000" | jq '.. |
objects | select(.data_type? | strings | startswith("timestamp"))
| .value'
```

— Flags, output fields, and a runnable example are all specified. v2026.11.24

02 Regex filtering of JSON output

Not checked yet NEW

80

The new `-f` flag filters JSON output with a regex: in nested mode it keeps matching nodes plus their ancestors (marked `matched`), while in flat mode it keeps only the matching nodes.

Isolate only timestamp nodes from a batch file of URLs using the JSON filter, keeping ancestors for context.

```
$ unfurl -t json -f timestamp urls.txt
```

— Exact flag behavior per output mode and a runnable example given. v2026.11.24

03 **Timeflake ID parser** Not checked yet NEW 70

Unfurl adds a parser that recognises Timeflake 128-bit IDs in 32-hex, 22-base62, or dashed UUID form, identifying them by length and a plausible embedded timestamp from 2020 onward.

— Mechanism and formats detailed, but no example command shown. v2026.11.24

THINNER COVERAGE BELOW

04 **Notion and Kimi ID parsers** Not checked yet NEW 20

Unfurl also adds parsers to recognise Notion and Kimi IDs embedded in URLs, alongside the new Timeflake parser.

— Parsers are only named; no mechanism or example provided. v2026.11.24

05 **Short-link expansion safety limits** Not checked yet NEW 10

Unfurl adds safety limits when expanding short links, though no further mechanism or configuration detail is given.

— Mentioned only in the release summary with no detail. v2026.11.24

WAS THIS USEFUL?  

ALEAPP

Feedback ⓘ v2026.4.6 NOTES ↗ CODE ↗ RANK

ALEAPP (Android Logs Events And Protobuf Parser) is a digital forensics tool that parses and analyzes Android device artifacts for investigative purposes.

> ALEAPP's v2026.4.6 release adds Signal on-device backup decryption, dozens of new artifact parsers spanning Samsung, Google, messaging apps and system/app data, native Linux .deb/.rpm packages, and a round of report-rendering fixes.

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

01 **New artifact parsers** Not checked yet NEW 90

Adds new artifact parsers across platforms: Samsung Messages, visual voicemail, Phone RTT call text, Quick Share transfers and link shares, Find devices and locations, Reminder, Voice Recorder, Gallery search and operation history, Clock alarms and world clocks, Calculator history, Internet bookmarks, saved pages, tabs, search history, quick access and web notifications, Health weight, steps and medications, Digital Wellbeing timers, context engine, Secure Folder file operations, vold external storage events, Customization Service screen state, Cloud Clock Offset, and encrypted Health and Customization Service stores; Google Find My Device locations and devices, GMS semantic location with activity and Wi-Fi, Contact Keys, Google Accounts from `accounts_ce.db`, the Play Services SMS index, the advertising ID, AdServices attribution and consent, Calculator history, Chrome Accounts and Reading List, and Chromium Favicons; WhatsApp status updates and viewers, Business tools, local and Google Drive backups, and the default disappearing messages timer; Signal conversation view, account, image cache, and on-device backups, decryptable by supplying the account's 64-character recovery key with `--signal_recovery_key_file` or `--signal_recovery_key_env` (without a key the backup folder is listed but not decrypted); Threema account, contacts, messages and media; Jami; Zangi calls and groups; newer Wire databases when the extraction carries the keys; Telegram downloaded media linked from `file_to_path.db` with Saved Messages chat labelled; LINE messages without a contacts table; App Ops access history on Android 16, Telephony Analytics, URI permission grants, companion device associations, enabled accessibility services, Android Keystore entries, wallpapers, Photo Picker synthetic files, SMS device clock gap, bug report shortcuts, Beats devices, and battery usage on v10 and v11 databases for every user; and Snapchat Conversation Turn Events, Last Saved Location and Client Notifications, TikTok search history and play records, Spotify stores, Stock Android Browser and Email, KDL artifacts, Smart Select fragments, the Protected Notes app, and a LAVA conversation view for ChatGPT.

Parse a full-filesystem extraction that includes Signal on-device backups, supplying the recovery key so messages, attachments, recipients and chats are decrypted during the run.

```
$ aleapp -t zip -i /cases/device_extraction.zip -o /cases/output/
--signal_recovery_key_file /cases/signal_recovery.key
```

— Names exact flags and shows a runnable command Lists many concrete artifacts but no usage mechanism Names concrete Google artifacts including a specific database path Names many messaging apps and a specific linking database Names specific system artifacts and OS version scope Names several apps and artifact types

v2026.4.6

02 **Lx01 and ASIF support for raw disk images** Not checked yet

IMPROVED 65

Reading disk images with `-t raw` now also accepts Lx01 and ASIF logical evidence formats, in addition to existing raw image support.

Process a raw disk image that is an Lx01 logical evidence container — now supported directly without mounting.

```
$ aleapp -t raw -i /cases/evidence.Lx01 -o /cases/output/
```

— Names the flag and new formats with an example command v2026.4.6

03 **Report rendering and formatting improvements** Not checked yet

IMPROVED 60

LAVA report output header types `number`, `latitude` and `longitude` are now stored as NUMERIC, case data is recorded in the LAVA manifest, and report cells break on every line ending; HEIC media now renders in reports as a JPEG copy generated with heicreader instead of pillow-heif, removing all GPL and LGPL code from packaged builds; HTML email bodies now render as readable text, K-9 plain text parts are decoded, and DeepSeek message text is stored without HTML escapes.

— Names exact header types changed but no usage path Explains mechanism and licensing reason, no user action Names affected apps but no configuration surface

v2026.4.6

THINNER COVERAGE BELOW

04 **Linux .deb and .rpm packages** Not checked yet NEW 50

Linux packages now include `.deb` and `.rpm` for x64 and ARM alongside the existing AppImages.

— Names exact package formats and architectures v2026.4.6

05 **walStrings artifact removed** Not checked yet BREAKING 50

The `walStrings` artifact has been removed; it listed ASCII text runs from `-wal` and `-journal` files and was already off by default.

— Names the removed artifact and its prior behaviour v2026.4.6

---

06 **netstats output enhancements** Not checked yet IMPROVED

40

netstats output now shows identity fields, each row's collection and source file, and continues processing when one input is damaged.

— Describes behaviour change but no concrete fields named v2026.4.6

└──▷ **BREAKING ON UPGRADE**

! The `walStrings` artifact has been removed; it listed ASCII text runs from `-wal` and `-journal` files and was already off by default.

WAS THIS USEFUL? 👍 👎

iLEAPP

Feedback ⓘ v2026.4.7 NOTES ↗ CODE ↗ RANK

iLEAPP is an open-source iOS forensic artifact parsing tool that extracts and analyzes data from iPhone backups and filesystem images.

> iLEAPP adds native parsing of raw Lx01/ASIF disk images, decrypts the Find My encrypted location cache using the extraction's keychain key, and ships dozens of new and reworked artifacts spanning Biome, knowledgeC, Health, PowerLog, WhatsApp and numerous third-party apps, alongside LAVA output schema improvements and a reworked GUI.

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

01 Raw Lx01/ASIF disk image parsing

Not checked yet NEW

80

iLEAPP can now read disk images in Lx01 and ASIF formats directly via `-t raw` on the CLI, without a separate conversion step.

Parse a raw disk image in Lx01 or ASIF format, now supported directly without conversion.

```
$ ileapp -t raw -i /path/to/image.lx01 -o /path/to/output/
```

— runnable CLI command with named flag and formats v2026.4.7

02 New Biome artifacts and streams

Not checked yet NEW

60

Adds new Biome artifacts and streams: Motion Activity, Location Semantic, Local Authentication Evaluations, Photos Delete/Share/Engagement, Widgets Viewed, Focus Driving, Inferred Mode, Sleep Mode, Find My Location Change, NFC Tag, nearby devices, App Extension Usage, Phone Indexed Calls, Siri Execution, Notification Delivery, and Screenshots; Apple Mail `mailIntelligence` Biome stream events (app launch, resume, background, message fetch, view, read, sent, moved, link clicked); and Biome streams for `FindMy.LocationChange`, `Device.Wireless.NFCTag`, `Device.Wireless.DefaultPairedNearby`, and `Device.Wireless.BluetoothNearbyDevice`.

— exhaustive named artifact list but no usage example v2026.4.7

THINNER COVERAGE BELOW

03 LAVA output schema enhancements

Not checked yet IMPROVED

55

Adds `number`, `latitude`, and `longitude` header types stored as NUMERIC in LAVA output, adds hash columns typed as hashes, and records architecture in run info.

– named schema fields but no example output v2026.4.7

04 **New third-party app artifacts** Not checked yet NEW 55

Adds new app artifacts for Citizen, Dust, Locket, Wire (accounts, contacts, media, conversation view), PayPal, Zelle, AnyDesk, OfferUp, Craigslist, Google Measurement (Firebase Analytics), Maps Sync typed and sent queries, Shortcuts, Jami, Calculator history, FaceTime Links, Weather, Configuration Profiles, Cloud Configuration Details, Background URL Session Tasks, Media Playback Event Streams, Omnibox Shortcuts, Google app tabs, and WebKit website data per app; a Zillow Recent Searches artifact parsed from `ZRECENTSEARCH` in Zillow.sqlite; and three Weather app artifacts: Last Fetch by Coordinate, Current Location, and Reverse Geocode Results.

– long named app/field list, no usage example v2026.4.7

05 **PowerLog artifact additions** Not checked yet NEW 55

Adds PowerLog artifacts for App Run Time per aggregation window, In Call Service, Lightning Connector, Network Usage, Time Zone, locationd Motion State History parsed from `cache_encryptedC.db`, and netusage Connections fields.

– named source file but no command to run it v2026.4.7

06 **Find My encrypted cache decryption** Not checked yet NEW 50

Adds a Find My encrypted cache decryption capability that uses the keychain key obtained from the extraction to decrypt the cache.

– mechanism named but no command or output shown v2026.4.7

07 **knowledgeC artifact expansion** Not checked yet IMPROVED 45

Adds knowledgeC artifacts for Application In Focus, notifications, app activity, web usage, and Location Activity payloads decoded by named field, and reworks six existing knowledgeC artifacts.

– named artifact types but no mechanism detail v2026.4.7

08 **New GUI window and config-driven versioning** Not checked yet NEW 45

Adds a new GUI window; the tool name and version shown now come from `data/leapp_config.json` instead of being hardcoded.

– named config file gives clear starting point v2026.4.7

09 **ileappGUI executable removed** Not checked yet

BREAKING

45

The separate `ileappGUI` executable is removed; running `ileapp` without arguments now opens the GUI window instead.

— clear migration note on invoking the GUI v2026.4.7

---

10 **Health artifact additions and rework** Not checked yet IMPROVED

40

Adds a Health Flights Climbed artifact and reworks five Health artifacts to use HealthKit flight and distance type codes.

— named artifact and codes but no usage path v2026.4.7

---

11 **WhatsApp message receipts and LID data** Not checked yet NEW

40

Adds a WhatsApp Message Receipts artifact and WhatsApp LID phone numbers, group members, and linked group member per message.

— named fields but no mechanism or example v2026.4.7

└──▷ **BREAKING ON UPGRADE**

! The `ileappGUI` executable is removed; `ileapp` without arguments now opens the window instead.

WAS THIS USEFUL? 👍 👎

ElcomSoft

Feedback ⓘ announced · Nov 23, 2026 BLOG ↗ RANK

ElcomSoft is a try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

> ElcomSoft's Quick Triage tool gains native mounting of forensic and VM disk images across many formats with a new per-volume mounting layer, plus an Image Management module, archive mounting, drive-letter reconstruction, and reorganized smart folders for artifact triage.

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

01 **Native mounting of forensic and VM disk images** Not checked yet

NEW 85

Natively mounts forensic and VM disk images in `RAW/DD`, `E01`, `AFF4`, `VHD/VHDX`, `VDI`, `VMDK`, `Parallels`, and `QCOW/QCOW2/QED` formats, with `NTFS`, `FAT`, and `exFAT` partitions on `MBR` and `GPT` disks. Partitions mount as folders using a new per-volume MetaFS layer instead of drive-letter assignment, avoiding drive-letter exhaustion on images with many partitions.

Triage a forensic E01 image from a seized machine without assigning drive letters, then browse its partitions directly in Explorer.

> 📍 In Elcomsoft Quick Triage, go to File › Add Evidence › Disk Image, select your .E01 file, and let EQT ingest it. Open the Image Management module to view partition metadata, then click `Mount` on the target partition to open it as a folder in Explorer.

— Names formats and mounting mechanism, with steps to mount launch-20261123-728f1e10

02 **Image Management module for mounted volumes** Not checked yet NEW

65

The new Image Management module lets examiners view partition metadata, batch mount or unmount volumes, open partitions in Explorer, and relink moved image files with Image ID verification.

Triage a forensic E01 image from a seized machine without assigning drive letters, then browse its partitions directly in Explorer.

> 📍 In Elcomsoft Quick Triage, go to File › Add Evidence › Disk Image, select your .E01 file, and let EQT ingest it. Open the Image Management module to view partition metadata, then click `Mount` on the target partition to open it as a folder in Explorer.

— Named module demonstrated with a mount-and-browse workflow launch-20261123-728f1e10

THINNER COVERAGE BELOW

03 **Automatic OS detection and drive-letter reconstruction**

Not checked yet NEW 55

Automatically detects the operating system and reconstructs original Windows drive-letter assignments when ingesting disk images, eliminating manual path translation for Registry hives, Jump Lists, LNK files, and download histories.

— Explains purpose but no concrete steps given launch-20261123-728f1e10

04 **Mount ZIP and TAR archives as virtual images** Not checked yet

NEW 45

Mount ZIP and TAR archives directly as virtual images without decompression, making the directory tree browsable in Windows Explorer and ready for triage parsing.

— Clear capability but no walkthrough or example given launch-20261123-728f1e10

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.16.15 NOTES ↗ CODE ↗ RANK

Netcap is a framework for secure and scalable analysis of network traffic, converting pcap files into typed audit records for efficient processing.

> Netcap reworked all 37 built-in detection rules to count per actor and retired three noisy rules (removing `configs/firewall-rules.yml`), while adding JA3/JA3S TLS fingerprinting, new SMB/HTTP/onion/UDP detections, an alert API that separates findings from inventory, and fixes for BPF-filtered and Token Ring capture decoding.

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

01 Alert API finding/inventory classification and fields

Not checked yet NEW 90

Alert views now support a `?kind=finding|inventory|all` query parameter, defaulting to `finding` to suppress inventory noise from the findings view, and the `AlertResponse` JSON now includes `kind`, `count`, `firstSeenNs`, `lastSeenNs`, `aggregateKey`, and `examples` fields.

Filter the alert API to show only actionable findings, excluding routine inventory observations like `SMB seen` or `Apache seen`.

```
$ curl 'http://127.0.0.1:8080/api/alerts?kind=finding'
```

Review only inventory-class alerts (services observed, cleartext protocols in use) separately from security findings.

```
$ curl 'http://127.0.0.1:8080/api/alerts?kind=inventory'
```

— Exact endpoint param and fields given with runnable curl examples. v0.16.15

02 Built-in detection rule rework and retirements Not checked yet

85

BREAKING

All 37 built-in threshold rules now count per source, source–destination pair, receiving client, victim, or distinct destination instead of across every host together, with thresholds above 256 counting in time buckets with per-rule state budgets; DNS Beaconing detection now matches queries case-insensitively via `fold_case`, reducing false positives from name-case variation; and the Port Scan High Port, LDAP Replication

Activity, and Missing Security Headers rules are retired, removing `configs/firewall-rules.yml` and dropping saved copies of retired rules on load.

— Names all retired rules, removed config file, and matching mechanism. v0.16.14

03 **BPF-filtered capture decoding fix** Not checked yet IMPROVED

85

BPF-filtered captures now use the bounded reader and write `capture-input.json` with a `bpfExcludedPackets` field recording how many packets were excluded by the filter; records whose `caplen` exceeds `origlen` are now decoded and tracked as `origLenRaisedPackets`, so captures such as `sip.pcap` process all packets instead of none.

— Named output file, fields, and a concrete before/after example given. v0.16.14

04 **Alert group resolution and oversized record handling**

Not checked yet IMPROVED

75

Alert group resolution now records each group once in `resolved_alerts.json` version 2, and matched records above 256 KiB are replaced by their identity and hash; older alerts above 1 MiB are skipped and counted as `skippedOversizedAlerts`.

— Names file, version, size thresholds and counter field. v0.16.14

THINNER COVERAGE BELOW

05 **Artifact classification and extraction reporting** Not checked yet

IMPROVED

55

Artifact classification now covers HTML, JavaScript, Flash, and WOFF files, and reports partial extractions and artifact-input loss separately from analysis failures.

— Names formats and distinguishes failure types, no exact fields. v0.16.15

06 **Showcase view capture summaries and limits** Not checked yet

IMPROVED

50

Showcase view now includes capture summaries with evidence links, bounded Files and Fingerprints views, visible UDP retention limits, and a detection-rule pack digest in service run records.

— Lists UI additions without exact navigation or values. v0.16.15

07 **Service mode capture scoping and defaults** Not checked yet

IMPROVED

45

Service mode now scopes capture selection per browser tab, with private uploads scoped to the visitor and shared showcase results shown as read-only; it also uses the same analysis defaults as the CLI and accepts partial captures without a capture manifest.

— Describes scoping and defaults but no config keys or UI paths. v0.16.15 v0.16.14

08 **New traffic detections for SMB, HTTP, onion, and UDP**

Not checked yet NEW 40

Netcap now detects SMB probe and backdoor responses, command-bearing HTTP requests, onion names in cleartext traffic, and UDP fan-out.

— Lists four detection types with no logic or thresholds given. v0.16.15

09 **WebUI alert list streaming and pagination** Not checked yet

IMPROVED 40

The WebUI alert list and grouped view now stream instead of loading the full audit, and the Alerts page pages correctly through large groups.

— Describes behavior change but no concrete limits or path. v0.16.14

10 **JA3/JA3S TLS fingerprint recording and filters** Not checked yet

NEW 35

Netcap now records JA3 and JA3S TLS fingerprints and adds fingerprint filters.

— Names fingerprint types but no filter syntax or UI path. v0.16.15

11 **Token Ring (DLT 6) decoding support** Not checked yet NEW

35

Netcap can now decode Token Ring (DLT 6) captures that were previously unreadable.

— Names the link-layer type but no decoding detail. v0.16.14

12 **Passive CVE findings require version match** Not checked yet

IMPROVED 30

Passive CVE findings now require an applicable version or version range before surfacing.

— States the behavior change with no example or threshold. v0.16.15

└──▷ **BREAKING ON UPGRADE**

! `configs/firewall-rules.yml` is removed; saved copies of retired rules (Port Scan High Port, LDAP Replication Activity, Missing Security Headers) are dropped on load.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Mend.io Mend Container

Feedback ⓘ

changelog · Nov 25, 2026 CHANGELOG ↗

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

NEW SUPPORTED PLATFORM Red Hat Hardened Images (Hummingbird)

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ

Exposure checks behind a large AWS WAF web ACL update again: NOTES ↗ RANK

Plerion is a cloud security platform that provides risk-based prioritization and remediation guidance across cloud infrastructure, helping teams identify and fix critical vulnerabilities efficiently.

> Plerion shipped a new Kubernetes workload public-exposure check spanning EKS, GKE, and AKS, source-based workflow matching for CDR findings, and a cluster of accuracy fixes to S3, RDS/EMR, and Macie-backed checks that will change pass/fail results on the next scan.

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

01 Kubernetes workload public-exposure check

Not checked yet NEW

85

New check `PLERION-K8S-188` detects publicly accessible Kubernetes workloads by tracing each workload's path from a public load balancer, Ingress, or Gateway to its pods across EKS, GKE, and AKS, accounting for network policies including Calico, Cilium, and Istio. It requires collector-manager v1.2.13 or later.

— Names check ID, mechanism and version requirement; no setup steps given

Exposure checks behind a large AWS WAF …

02 Source-based workflow matching for CDR findings

Not checked yet

NEW

85

Workflows can now match CDR findings by source — GuardDuty, CloudTrail, Macie, or IAM Access Analyzer — instead of maintaining check ID lists, so every new detection from that source is automatically included going forward.

Route all GuardDuty CDR findings to a notification channel without maintaining a list of check IDs — new detections are picked up automatically.

📍 In Workflows, create a new rule, set the Source filter to `GuardDuty`, and configure your desired action (e.g. send to Slack or PagerDuty). Save — all current and future GuardDuty findings will match.

— Exact UI steps included for routing findings without check IDs

Workflows for CDR findings, no check ID…

03 S3 sensitive-data checks align with Macie findings

Not checked yet

85

BREAKING

Checks `PLERION-AWS-857`, `PLERION-AWS-858`, `PLERION-AWS-859`, and `PLERION-AWS-860` now count a Macie `SensitiveData:S3Object/Multiple` finding toward each sensitive-data category it contains, so a bucket holding both credentials and card numbers fails both the credentials and financial checks. Each check now fails a bucket only when Macie holds an active finding of that check's own category and returns no result otherwise, so buckets never analysed by Macie no longer pass; buckets that previously passed based solely on a Multiple finding will now show as failed.

— Names all four check IDs and the exact pass/fail logic change

Macie checks see objects with more than…

04 **CloudTrail data-event checks require full-bucket selectors**

Not checked yet

75

BREAKING

S3 CloudTrail data event checks `PLERION-AWS-38`, `PLERION-AWS-39`, and `PLERION-AWS-854` now pass only when a CloudTrail selector covers the entire bucket; selectors scoped to a prefix (e.g. `arn:aws:s3:::my-bucket/logs/`) no longer pass the check.

— Names three check IDs and the exact selector scope rule

Workflows for CDR findings, no check ID…

05 **S3 access-logging check reads server access logs** Not checked yet

70

BREAKING

Check `PLERION-AWS-854` ('S3 bucket does not have server access logging enabled') now reads server access logging instead of CloudTrail data events, passing buckets that log to another bucket or to CloudWatch Logs. Buckets that previously passed based on CloudTrail data events alone will now fail on the next scan.

— Names check ID and before/after but no remediation steps

Exposure checks behind a large AWS WAF …

06 **RDS and EMR exposure checks evaluate all load balancers**

Not checked yet

60

BREAKING

Plerion's fact check for RDS and EMR now connects to every Network Load Balancer in front of a resource, not just one, returning Factual when any NLB answers and Unknown when a load balancer is unreadable and none respond. Resources behind two or more internet-facing NLBs that previously passed may now fail on the next scan.

— Explains mechanism and impact but no remediation path

Workflows for CDR findings, no check ID…

THINNER COVERAGE BELOW

07 Macie finding retention and ordering Not checked yet IMPROVED

45

Plerion now retains every Macie finding for objects whose results differ across scans, and surfaces the newest findings first when an account holds more than 10,000 findings.

— States behavior and a threshold but no further mechanism

Macie checks see objects with more than…

└──▷ BREAKING ON UPGRADE

! Buckets that previously passed `PLERION-AWS-854` based on CloudTrail data events alone will now fail on the next scan, because the check now evaluates server access logging instead.

! Findings for RDS and EMR resources behind two or more internet-facing Network Load Balancers that previously passed may now fail on the next scan, as the fact check now evaluates all NLBs rather than just one.

! S3 CloudTrail data event checks (`PLERION-AWS-38` , `PLERION-AWS-39` , `PLERION-AWS-854`) now pass only when a CloudTrail selector covers the entire bucket; selectors scoped to a prefix (e.g. `arn:aws:s3:::my-bucket/logs/`) no longer pass the check.

! Buckets that previously passed `PLERION-AWS-857` , `PLERION-AWS-858` , `PLERION-AWS-859` , or `PLERION-AWS-860` because their only Macie finding was `SensitiveData:S3Object/Multiple` will now show as failed; expect passed findings to disappear and new failed findings to appear on the next scan.

! Buckets with no Macie finding no longer pass any of the four S3 sensitive-data checks; they now return no result instead of a pass, because an absence of findings may mean Macie never analysed the bucket or findings reside in an unreadable Macie administrator account.

WAS THIS USEFUL? 👍 👎

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Nov 23, 2026 BLOG ↗ RANK

Rapid7 InsightCloudSec is a cloud-native security platform that continuously monitors, detects, and automatically remediates misconfigurations and compliance risks across multi-cloud environments.

> InsightCloudSec added new compliance checks for AWS database services and a database change supporting custom plugins in self-hosted deployments.

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

01 **CIS AWS Database Services Benchmark insights for ElastiCache, RDS, Neptune**

Not checked yet NEW 45

InsightCloudSec adds Insights and Query Filters covering Amazon ElastiCache, RDS, and Neptune aligned to the CIS AWS Database Services Benchmark v2.0.0.

— Names services and benchmark version but no specific filter names 2026-11-23-078b63fd

02 **Database change for custom plugins in self-hosted deployments**

Not checked yet IMPROVED 20

Self-hosted deployments now include a database change intended to support custom plugins.

— Bare description with no mechanism or specifics given 2026-11-23-078b63fd

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Turbot Steampipe

Feedback ⓘ v2.5.3 NOTES ↗ CODE ↗ RANK

Steampipe is an open-source tool that lets you query cloud infrastructure, APIs, and security data using SQL.

> Steampipe v2.5.3 speeds up aggregator connection startup and tightens validation of database connection URLs, alongside smaller fixes to plugin timeouts and parallelism limits.

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

01 Database URL password encoding requirement

Not checked yet

83

BREAKING

A remote database URL whose password contains an unencoded `/`, `?`, `#`, or `@` is now refused; such characters must be percent-encoded, e.g. write `@` as `%40`. A raw `@` that previously worked in v2.5.2 is now rejected, an `@` after the host is refused even as `%40`, and `#` is also disallowed — pass an `sslrootcert` path containing `@` via the `PGSSLROOTCERT` environment variable instead.

— Names exact characters, env var, and migration workaround v2.5.3

THINNER COVERAGE BELOW

02 Faster aggregator connection startup

Not checked yet IMPROVED

50

Aggregator connections with two or more child connections are no longer re-imported on every service start when nothing has changed, reducing startup time on databases with large aggregators.

— Explains mechanism and benefit but no config detail v2.5.3

03 Clearer STEAMPIPE_MAX_PARALLEL error messages

Not checked yet IMPROVED

35

Steampipe now surfaces clearer errors when the `STEAMPIPE_MAX_PARALLEL` value is set too high.

— Names config var but no detail on new message or threshold v2.5.3

04 5-minute plugin timeout

Not checked yet NEW

25

Steampipe now applies a 5-minute timeout for plugins.

— Mentioned only in summary, no mechanism given v2.5.3

└──▷ **BREAKING ON UPGRADE**

! A remote database URL whose password contains an unencoded `/`, `?`, `#`, or `@` is now refused — write `@` as `%40`. A raw `@` that worked in v2.5.2 is rejected. An `@` after the host is refused even as `%40`, and so is a `#`; pass an `sslrootcert` path containing `@` via `PGSSLROOTCERT` instead.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ **SOAR**

Torq

Feedback ⓘ changelog · Nov 25, 2026 CHANGELOG ↗ RANK

Transform the enterprise SOC with Torq — combining agentic SecOps and Hyperautomation to triage, investigate, and respond to threats faster than ever.

> Torq added workflow steps and public API support for managing Cases Dashboards and Scheduled Reports across workspaces.

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

01 **Cases Dashboards management via workflow and API** Not checked yet

NEW 55

New workflow steps and Torq public API endpoints let users manage Cases Dashboards, including exporting a dashboard as a template and importing it into another workspace.

— Names mechanism (export/import template) but no exact API endpoint

changelog-20261125-c4006d23

02 **Scheduled Reports management via workflow and API** Not checked yet

NEW 44

New workflow steps and the Torq public API enable managing Scheduled Reports, allowing report scheduling across workspaces from a single workflow.

— Describes capability but no specific endpoint or flag names changelog-20261125-c4006d23

WAS THIS USEFUL? 👍 👎

Tines

Feedback ⓘ Restrict access to a specific case (2026-11-24)

BLOG ↗ RANK

Tines is an apps, agents and automations, in the complete environment for enterprise AI.

> Tines added per-case access restriction, letting teams limit visibility of sensitive cases to a selected set of users.

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

01 **Restrict case access to selected users** Not checked yet NEW

40

Case access can now be restricted to a selected set of users within a team, useful for sensitive investigations, HR incidents, or cases that should not be shared team-wide.



— Describes purpose but no config path or mechanism named

Restrict access to a specific case (202…

WAS THIS USEFUL? 👍 👎

◆ **Endpoint Security (EDR/XDR)**

Sophos Central

Feedback ⓘ changelog · Nov 25, 2026

CHANGELOG ↗

Sophos Central is a cloud-based management platform that enables administrators to centrally manage and monitor Sophos security products across endpoints, networks, and servers.

NEW DETECTION Image Analysis for Sophos Email Plus: OCR

WAS THIS USEFUL? 👍 👎

Huntress

Feedback ⓘ

Choose how Huntress notifications are prioritized in Syncro NOTES ↗ RANK

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

> Huntress added severity-based ticket priority mapping to its Syncro integration, letting teams align Huntress alerts with their existing triage workflow.

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

01 **Severity-to-priority ticket mapping in Syncro integration**

Not checked yet NEW 75

The Syncro integration gains Ticket Priority Mappings on the Map Tickets tab, allowing Huntress notification severities (Critical, High, and Low) to be mapped to corresponding Syncro ticket priorities.

Align Huntress-generated Syncro tickets with your team's triage workflow by mapping each Huntress severity to the correct ticket priority.

> 📍 In the Syncro integration, go to the Map Tickets tab › Ticket Priority Mappings, then set the desired Syncro ticket priority for each of Critical, High, and Low Huntress notification severities, and save.

— Names exact UI path and mapping values; no mechanism beyond config

Choose how Huntress notifications are p…

WAS THIS USEFUL? 👍 👎

XFA

Feedback ⓘ See which apps your people sign in to NOTES ↗ RANK

XFA is a developer tool that enables building fast, interactive forms and data collection interfaces for web applications.

> XFA introduced a beta App usage view that lets teams see which apps employees sign in to with their work accounts.

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

01 **App usage (beta) view** Not checked yet NEW 40

A new App usage BETA view shows which apps people sign in to with their work account, browsable per app and per person. In XFA, navigate to App usage BETA and select a person or an app to see sign-in activity for that account.

Review which apps your workforce is signing in to with work accounts, to identify unsanctioned or risky SaaS usage.

> 📍 In XFA, navigate to App usage BETA and select a person or an app to see sign-in activity for that account.

— Names UI location but no mechanism, data source, or limits.

See which apps your people sign in to

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.261124.0 NOTES ↗ RANK

OpenCTI is an open-source threat intelligence platform that helps organizations manage, analyze, and share structured cyber threat intelligence data.

> OpenCTI's 7.261124.0 release adds a GraphQL mutation for attaching staged sync files, admin controls for outbound-request blocking and session lifetime, dashboard/feed export-import, platform-wide search, and a set of security hardening defaults for credentials and SMTP.

└──▷ **WHAT SHIPPED · 9 FEATURES** most completely described first ? what's the number?

01 **importPushRef GraphQL mutation for sync feeds** Not checked yet

NEW 75

A new `importPushRef` GraphQL mutation attaches staged sync files to an existing sync feed, letting admins finalize bulk-imported files against a feed via the API.

Attach a staged sync file to an existing sync feed via the GraphQL API after a bulk import.

```
$ curl -X POST https://<opencti-host>/graphql \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"query": "mutation { importPushRef(id: \"<sync-feed-id>\",
file: \"<staged-file-ref>\") { id } }"}'
```

— Named mutation with a runnable curl example 7.261124.0

02 **Security hardening defaults** Not checked yet NEW 60

New security hardening controls include an admin switch to block outbound requests to internal networks (reducing SSRF-style exposure), a requirement for TLS on SMTP connections when credentials are configured, startup validation of the admin password against the password policy, and refusal of default credentials for the admin account and backing services.

— No named flag or exact UI path given Names three concrete startup validation checks

7.261124.0

03 **Feed and dashboard export/import via config bundle wizard**

Not checked yet NEW 60

The configuration bundle wizard now supports export and import of feeds in `RSS`, `CSV`, `JSON`, and `TAXII` formats, as well as dashboards.

— Names formats and the wizard entry point 7.261124.0

THINNER COVERAGE BELOW

04 **Source code asset category and purl attribute** Not checked yet

NEW

45

The system taxonomy gains a `source code` asset category, and the Software observable gains a single-valued `purl` attribute.

— Names the taxonomy and attribute added 7.261124.0

05 **Global platform page search** Not checked yet NEW

40

Global search now covers platform pages and sections, not just data objects, letting users search for UI destinations directly.

— Describes scope but no exact navigation given 7.261124.0

06 **Maximum session lifetime enforcement** Not checked yet NEW

40

Admins can now end user sessions after a maximum lifetime is reached even if the user remains active, adding a hard session expiry on top of idle timeouts.

— Behaviour described but no config key named 7.261124.0

07 **opencrq environment declaration and product switcher**

Not checked yet NEW

35

`opencrq` can now be declared by environment variable and is surfaced in the product switcher.

— Names the variable but gives no further detail 7.261124.0

08 **Functional date selection in dashboard table widget**

Not checked yet NEW

30

The dashboard table widget now supports selecting a functional date for its data.

— Bare description of a single widget change 7.261124.0

09 **Pastel colour user preference** Not checked yet NEW

30

Users can now select pastel colours as a display preference, with pastel remaining the default colour scheme.

— Minor cosmetic preference, thinly described 7.261124.0

WAS THIS USEFUL? 👍 👎

IT-Harvest

Feedback ⓘ announced · Nov 24, 2026 BLOG ↗ RANK

Data-driven analyst firm covering the entire cybersecurity industry, with 4,000+ vendors and 11,000+ products in one living research platform (the IT-Harvest Dashboard).

> IT-Harvest launched CyberExchange, a new product-first research platform that replaces the IT-Harvest Dashboard and adds AI-agent access, framework-based stack-gap analysis, and portfolio alerting.

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

01 **Framework coverage mapping and stack-gap comparison**

Not checked yet NEW INTEGRATION NEW 75

Map your existing tool stack to NIST CSF 2.0, MITRE ATT&CK, MITRE D3FEND, and CIS Controls, then compare coverage before and after adding a candidate product to see its marginal contribution.

Evaluate a candidate product's marginal value before purchasing by mapping your current stack and comparing MITRE ATT&CK coverage with and without it.

📍 In CyberExchange, go to See what your stack is missing, enter your existing products, select a candidate, and review the before/after coverage comparison across MITRE ATT&CK, NIST CSF 2.0, MITRE D3FEND, and CIS Controls.

– Exact navigation path and mechanism for before/after comparison given.

launch-20261123-8cf4120e

02 **Slash commands in IQ chatbot** Not checked yet NEW 70

Use `/product`, `/capability`, or `/mitre` slash commands in the CyberExchange chatbot (IQ) for precise product lookups beyond plain-language search.

– Exact command syntax a reader can type directly. launch-20261123-8cf4120e

03 **MCP server access for AI agents** Not checked yet NEW 63

Connect AI agents to IT-Harvest vendor and product data via Model Context Protocol (MCP) servers, or download Markdown files for offline agent use.

Let an AI agent query CyberExchange vendor and product data directly without manual export, using the MCP server integration.

📍 Pick an MCP server at cyberexchange.ai, configure your agent to point at it, then issue natural-language queries against IT-Harvest's vendor and product dataset.

– Named protocol and two access modes, but no exact endpoint. launch-20261123-8cf4120e

04 **CyberExchange platform replaces IT-Harvest Dashboard**

Not checked yet

60

BREAKING

IT-Harvest adds CyberExchange as a full replacement for the IT-Harvest Dashboard, now covering 5,297 cybersecurity vendors, 15,233 products, and more than 20,000 open-source tools under a product-first lens. Existing Dashboard subscribers were transitioned automatically.

— Scope and scale given, but no migration steps. launch-20261123-8cf4120e

THINNER COVERAGE BELOW

05 **Multi-dimension product and vendor search** Not checked yet NEW

55

Search the new CyberExchange platform by MITRE ATT&CK technique, MITRE mitigation, NIST Cybersecurity Framework requirement, product name, capability, control, asset, attack method, or country.

— Many named search dimensions, but no query syntax shown. launch-20261123-8cf4120e

06 **Portfolio alerts for tracked vendors** Not checked yet NEW

45

Set up custom portfolio alerts to track acquisitions, headcount changes, funding rounds, and new products for shortlisted or in-stack vendors.

— Named alert categories but no setup steps given. launch-20261123-8cf4120e

└──▷ BREAKING ON UPGRADE

! The IT-Harvest Dashboard is replaced by CyberExchange; existing Dashboard subscribers were transitioned automatically.

announced · Nov 24, 2026 BLOG ↗

IT-Harvest adds the Open Source Security Center to `CyberExchange.ai` with facet search, preset Lenses, and per-tool Health, Continuity, Posture and Adoption scoring.

Search the catalog with facet queries such as `lang:rust`, `scorecard:>7`, `pushed:<30d`, and `has:signed-releases` to filter tools by language, OpenSSF score, recency, and release hygiene. Not checked yet

Browse tools through 10 preset Lenses including Maintained and permissive, Packaged for Kali, Hardened releases, Single-maintainer risk, and Mapped to MITRE. Not checked yet

> Each tool card exposes four 0-100 scores — Health, Continuity, Posture, and Adoption — plus OpenSSF Scorecard, open advisories, stars, forks, contributors, and last-commit date. Not checked yet
>
> Filter tools by platform, install method, output format, package registry, maturity, and GitHub topics. Not checked yet

---

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

rsigma

Feedback ⓘ v0.28.0 NOTES ↗ RANK

> RSigma v0.28.0 adds multi-tenant daemon support with per-tenant rulesets, threat-intel-driven detections from STIX/TAXII feeds, new dns and mmdb enrichers, and batched sink delivery that cuts 100,000-event NATS JetStream and OTLP delivery times by roughly 20x, alongside new OCSF/UDM output formats and a string of breaking changes to config validation, snapshots, and the library API.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 **Batched sink delivery performance** Not checked yet IMPROVED

100

Batch sink delivery now takes 100,000 events to NATS JetStream from 12.21 s to 0.65 s and to OTLP from 16.65 s to 0.59 s. Batch sizing is controlled with `--sink-batch-max` and flushing with `--sink-batch-flush-ms` (and `daemon.output.batch_flush_ms`); a zero value for `--buffer-size`, `--batch-size`, or `--sink-batch-max` is now rejected, and the `--sink-batch-flush-ms` default changes from 50 to 0.

— Exact benchmarks, flags and default change, fully actionable v0.28.0

02 **Multi-tenant daemon support** Not checked yet NEW

90

Adds multi-tenant daemon support with correlation, dedup, incident, silence, risk, and disposition state keyed per tenant; `daemon.tenancy.quotas` caps each tenant; API routes are scoped to the token's tenant; and sinks route by `{tenant}`. Per-tenant rule curation and pinning is configured via `rulesets:` in config, with a new `rsigma ruleset` command group and a `--ruleset` flag added to the `rule`, `eval`, `convert`, and `hunt` commands.

— Names config keys, commands and flags but gives no usage example v0.28.0

03 **DNS and MaxMind (mmdb) enrichers** Not checked yet NEW

83

Adds new `dns` and `mmdb` enrichers: `dns` expands DNS fields at detection time, and `mmdb` adds MaxMind database lookups. Setting `mmap: true` cuts idle memory from 113 MB to 21.5 MB on a 93 MB database.

— Concrete memory figures and config key but no CLI walkthrough v0.28.0

04 **Daemon CLI threading and stdin batch controls** Not checked yet

NEW

83

Adds daemon CLI controls `--worker-threads`, `--detect-threads`, and `--detect-inflight` for sizing thread pools on high-volume daemons, and `engine daemon --input stdin` to run one batch from stdin through to exit for serverless and cron callers. `--input stdin` now exits `3` when stored state cannot be restored and `4` when a run does not complete.

— Exact flags and exit codes given, ready to use `v0.28.0`

---

05 **Threat-intel-driven detections from STIX/TAXII** Not checked yet

NEW

78

Turns TAXII 2.1 collections and STIX bundles into dynamic source values for live detections, pushes indicators over the API with `watchlist` sources to rebuild referencing rules in real time, and drafts Sigma rules from STIX indicators with `rule draft --from-stix`. STIX ids outside the STIX 2.1 type-name rules no longer parse.

— Names API sources, flag and validation change, no example run `v0.28.0`

---

06 **Enrichment miss handling** Not checked yet BREAKING

65

Enrichment lookups that find nothing now trigger explicit `on_miss` handling, and the `http` enricher gains `default` and `miss_status` fields. An enricher configured with `default: null` now injects `null` on a miss instead of the previous behavior.

— Names fields and behavior change but no example `v0.28.0`

---

07 **Library API additions for Rust consumers** Not checked yet BREAKING

65

In the library API, `EnricherConfig` and `EnrichersConfigError` are now `#[non_exhaustive]`, `EnrichErrorKind` gains `Miss`, `CorrelationBody` gains `event_sources`, the tenancy types gain tenant fields, `RuleHeader` gains `ruleset`, `CorrelationSnapshot` gains `rulesets`, `SourceType` and `DataFormat` gain variants, `DraftError` and `StixIdError` gain variants, `TickOutput` and `RiskOutput` gain `finding_links`, and Fibratus `sanitize_field` is replaced by `field_ident`.

— Enumerates every changed API type verbatim `v0.28.0`

---

08 **New output formats: OCSF and UDM events** Not checked yet NEW

60

Adds new output format support: writing OCSF findings from eval, backtest, sinks, and webhooks, and canonicalizing Google SecOps UDM events with the new `udm_events.yml`.

— Names formats and file but no example output `v0.28.0`

---

09 **Strict config file validation** Not checked yet

BREAKING

60

Unknown keys in `rsigma.yaml` now stop every command with exit code `3`; unknown or inapplicable keys in enricher, webhook, risk, and source files fail to load.

— Exact exit code and affected files named v0.28.0

THINNER COVERAGE BELOW

10 **Rule-to-query conversion manifest** Not checked yet NEW

55

Convert groups of rules into queries plus a per-rule manifest using the `convert:` manifest with `rsigma backend build`.

— Names the command but no mechanism or example shown v0.28.0

11 **OTLP rejected records metric** Not checked yet NEW

55

Tracks OTLP rejected log records on the new `rsigma_sink_otlp_rejected_records_total{sink}` Prometheus counter.

— Names the metric but no usage context v0.28.0

12 **Correlation event lineage tracking** Not checked yet IMPROVED

50

Carries contributing `events` or `event_refs` and `event_sources` through chained correlations, naming the rule and child correlation behind each event.

— Names fields but no sample output shown v0.28.0

13 **Snapshot format version 2** Not checked yet BREAKING

35

Correlation, alert-pipeline, risk, and disposition snapshots move to version 2, which older releases cannot read.

— Names affected state but no migration path given v0.28.0

14 **Field identifier validation for LynxDB and Fibratus backends**

Not checked yet BREAKING

30

LynxDB and Fibratus conversion now rejects field names outside their identifier syntax.

— Names backends but not the identifier syntax itself v0.28.0

└──▷ BREAKING ON UPGRADE

! Unknown keys in `rsigma.yaml` now stop every command with exit code `3`; unknown or inapplicable keys in enricher, webhook, risk, and source files fail to load.

- ! An enricher with `default: null` now injects `null` on a miss instead of the previous behavior.
- ! `--input stdin` exits `3` when stored state cannot be restored and `4` when a run does not complete.
- ! A zero value for `--buffer-size`, `--batch-size`, or `--sink-batch-max` is now rejected.
- ! LynxDB and Fibratus conversion now rejects field names outside their identifier syntax.
- ! STIX ids outside the STIX 2.1 type-name rules no longer parse.
- ! Correlation, alert-pipeline, risk, and disposition snapshots move to version 2, which older releases cannot read.
- ! The `--sink-batch-flush-ms` default changes from 50 to 0.
- ! In the library API, `EnricherConfig` and `EnrichersConfigError` are now `#[non_exhaustive]`, `EnrichErrorKind` gains `Miss`, `CorrelationBody` gains `event_sources`, the tenancy types gain tenant fields, `RuleHeader` gains `ruleset`, `CorrelationSnapshot` gains `rulesets`, `SourceType` and `DataFormat` gain variants, `DraftError` and `StixIdError` gain variants, `TickOutput` and `RiskOutput` gain `finding_links`, and Fibratus `sanitize_field` is replaced by `field_ident`.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Genie Migrate conversion to Lakeflow pipelines is generally available (2026-11-24)

BLOG ↗ RANK

Databricks is a unified data analytics platform that enables collaborative data engineering, machine learning, and analytics using Apache Spark in the cloud.

> Databricks shipped a one-time PostgreSQL-to-Lakebase migration path and expanded classic instance pool configuration with automatic AZ selection, spot-with-fallback, and post-creation editability.

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

01 Migrate PostgreSQL databases into Lakebase

Not checked yet NEW

70

A one-time migration copies roles, schema, and data from an existing PostgreSQL database into a chosen branch of a Lakebase project, with no separate enablement step required. It can be run through the Lakebase App UI (select `Migrate`, point at the source database, choose the target branch) or via the Postgres API and SDKs.

Kick off a compatibility assessment and one-time data migration from a PostgreSQL database into a Lakebase branch via the UI.

> 📍 In the Lakebase App, open your Lakebase project, select `Migrate`, point it at your source PostgreSQL database, choose the target branch, and confirm to start the assessment and copy of roles, schema, and data.

— Mechanism and UI steps given, but exact API/SDK calls not detailed.

Migrate a Postgres database to Lakebase…

THINNER COVERAGE BELOW

02 Classic instance pool configuration options

Not checked yet

IMPROVED

54

Classic instance pools now support auto-AZ selection across all instance types, so Databricks automatically picks the optimal availability zone. Pools also support spot with fallback, letting spot instances fall back to on-demand when spot capacity is unavailable. Additionally, an existing pool's availability zone, availability mode, and max spot price can now be changed after creation without recreating the pool.

— Describes three related changes but no config keys or commands given.

New configuration options for classic i…

NEW INTEGRATION Genie Migrate: T-SQL to Lakeflow pipelines (GA): SQL Server, Azure Synapse dedicated SQL pools, Spark Declarative Pipelines

NEW DETECTION Streaming detection rules: SQL predicate, signature, indicator of compromise (IOC), atomic-signal detections

> Genie Migrate conversion to Lakeflow pipelines is generally available (2026-11-24)
>
> BLOG ↗ STABLE
>
> Databricks adds GA T-SQL migration to serverless Lakeflow pipelines, with per-group catalog/schema assignment and dry-run validation.
>
> Convert T-SQL scripts from SQL Server and Azure Synapse dedicated SQL pools into Lakeflow pipelines by selecting the Spark Declarative Pipelines target framework in a T-SQL migration project, organizing files into groups with a catalog and schema per group, each of which becomes one serverless pipeline validated by a dry run.
>
> Not checked yet
>
> Enable the Spark Declarative Pipelines target framework by having a workspace admin turn on Migration Projects Experimental Features from the Previews page.
>
> Not checked yet

> Lakewatch streaming detection rules (Experimental) (2026-11-24) BLOG ↗
>
> STABLE
>
> Databricks Lakewatch adds streaming detection rules EXPERIMENTAL with managed compute — no SQL warehouse or schedule selection required.
>
> Databricks Lakewatch now supports streaming detection rules that evaluate each new row of a source table as it arrives and create a signal for every row matching a SQL predicate, enabling low-latency signature, indicator of compromise (IOC), and atomic-signal detections without selecting a SQL warehouse or schedule.
>
> Not checked yet

WAS THIS USEFUL? 👍 👎

Elastic Security Labs

Feedback ⓘ announced · Nov 23, 2026

BLOG ↗

Elastic Security Labs is a powered by Elasticsearch. Ready to deploy your data?

Elastic Security 9.6 adds AlertZero agentic triage watches, enhanced Entity Analytics attention signals, ES|QL S3 Data Federation, Universal Alerting (GA), and an AI SRE.

AlertZero TECHNICAL PREVIEW adds four agentic Watches — Triage, Hunt, Detection, and Forensics — to work SOC alerts continuously with configurable autonomy and model choice. Not checked yet

Entity Analytics now highlights six attention signals for changing-risk entities and resolves identities across data sources. Not checked yet

ES|QL Data Federation TECHNICAL PREVIEW queries Parquet, CSV, TSV, and NDJSON files in Amazon S3 without ingestion. Not checked yet

The `document_extraction` inference task type TECHNICAL PREVIEW converts PDFs and documents to markdown using the `jina-ocr-v1` model. Not checked yet

Universal Alerting is now generally available, storing every ES|QL-written alert as searchable data alongside Classic Alerting. Not checked yet

+6 more

NEW INTEGRATION Vercel integration: Supabase, Temporal

NEW SUPPORTED PLATFORM Ransomware protection for Linux via Elastic Defend

NEW SUPPORTED PLATFORM FIPS 140-3 support in Elastic Agent: GovCloud

WAS THIS USEFUL? 👍 👎

◆ AI SOC

Exaforce

Feedback ⓘ Your AI Agent Coverage Is in Your Own Hands

NOTES ↗

Exaforce runs AI Exabots that detect, triage, investigate and respond to threats across cloud, SaaS, identity, endpoint and code on a real-time knowledge graph.

Exaforce adds agentless AI agent endpoint visibility via a once-daily MDM Snapshot Collector script and a unified AI & Agentic Insights dashboard.

NEW INTEGRATION Deploy the Snapshot Collector as a once-daily MDM script through Iru (Kandji), Intune, or another MDM to gain read-only endpoint inventory of AI skills, MCP servers, developer plugins, npm packages, and browser extensions — no persistent daemon or service required. Not checked yet

NEW INTEGRATION Get prompt blocking through official AI inference hooks, starting with Anthropic and Cursor, without routing traffic through a proxy. Not checked yet

Monitor all integration statuses, AI agentic metrics, and risk signals from a new unified AI & Agentic Insights dashboard. Not checked yet

NEW INTEGRATION EDR telemetry integrations: CrowdStrike FDR, SentinelOne Cloud Funnel, Defender XDR

NEW INTEGRATION Code repository integrations: GitHub, GitLab

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

oxo

Feedback ⓘ v2.25.8 NOTES ↗ RANK

> oxo added a new message type for subdomain fingerprinting data exchange between agents.

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

01 **Subdomain fingerprint message selector** Not checked yet NEW

60

Adds the `v3.fingerprint.domain_name.subdomains` message selector, enabling agents to publish and consume subdomain fingerprinting results within the OXO agent bus.

— Names exact message selector but lacks usage example v2.25.8

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

NinjaOne

Feedback ⓘ announced · Nov 23, 2026 BLOG ↗

NinjaOne is an improve IT efficiency with an autonomous patch management solution. Automate updates and secure endpoints across your organization with NinjaOne.

NinjaOne 16.0 adds agentless ChromeOS management via a Google Workspace integration, with remote actions, multi-tenant support, and device API access.

NEW SUPPORTED PLATFORM Enable ChromeOS device discovery and sync without deploying an agent by installing the NinjaOne Google Workspace app under Administration > Apps > NinjaOne Google Workspace. Not checked yet

Run remote actions — reboot, set volume, clear user profiles, or factory reset (powerwash) — on individual ChromeOS devices from the device page or in bulk from the device grid. Not checked yet

Access ChromeOS device inventory through NinjaOne's existing device API endpoints, so external reports and integrations can include ChromeOS devices from day one. Not checked yet

Connect multiple Google Workspace tenants, each mapped to a NinjaOne organization and location, so MSPs and multi-tenant IT teams can report across all connections in one console. Not checked yet

Open a Chrome Remote Desktop session directly from the ChromeOS device page for in-depth troubleshooting. Not checked yet

WAS THIS USEFUL? 👍 👎

Sophos Mobile

Feedback ⓘ 2026.47Released on November 23, 2026

NOTES ↗ RANK

Sophos Mobile is a unified endpoint management (UEM) solution that helps businesses manage and secure mobile devices, apps, and data across multiple platforms.

> Sophos Mobile added a configurable source for the MDM app download during device enrollment and clarified management status messaging in Sophos Fusion.

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

01 **Configurable MDM app download source for enrollment**

Not checked yet NEW 75

Admins can select a Sophos Manage source setting on the QR code enrollment and Samsung Knox Mobile Enrollment (KME) configuration pages to control where devices download the Sophos Mobile Device Management app; it defaults to managed Google Play but can be overridden for environments where that is blocked or unreachable.

Override the MDM app download source for environments where managed Google Play is blocked or unreachable during enrollment.

> 📍 In Sophos Mobile, go to the QR code enrollment or Samsung Knox Mobile Enrollment (KME) configuration page, locate the Sophos Manage source setting, and select an alternative download location.

— Names exact setting and config pages with override steps

2026.47Released on November 23, 2026

THINNER COVERAGE BELOW

02 **Clarified management status tooltip in Sophos Fusion**

Not checked yet IMPROVED 37

A tooltip on the Management status field of a device's details page in Sophos Fusion now clarifies that the status reflects the device's management mode, not full device control.

— Names exact field but minor UI text change only 2026.47Released on November 23, 2026

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Knocknoc

Feedback ⓘ changelog · Nov 25, 2026 CHANGELOG ↗ RANK

Knocknoc is a network access control tool that dynamically opens firewall rules to authenticated users, granting temporary IP-based access to protected services.

> Knocknoc shipped automatic Agent failover for continuous connectivity and added identity-based egress control for bastion hosts.

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

01 **Automatic failover for Agents** Not checked yet NEW 28

Knocknoc now supports automatic failover for Agents, keeping connectivity available when an Agent goes down.

— Describes benefit but no mechanism or config detail changelog-20261125-98f984c3

WAS THIS USEFUL? 👍 👎

DNSFilter Feedback ⓘ

11/24 November 24 – Web Application and API Update (2026-11-24) BLOG ↗ RANK

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

> DNSFilter's November update headlines a new clientless DoH-based filtering option alongside a revamped Audit Logs page, a new root CA, richer unblock-queue filtering, and an API enhancement for Roaming Clients.

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

01 Clientless DNS filtering via DoH address Not checked yet NEW

80

The new Clientless DNS Filtering feature under `Deployments > Clientless` lets routers, gateways, and any DoH-capable device apply DNSFilter policy through a unique DoH address, without needing a Roaming Client, Relay, or static IP. It is available on Plus plan and above; devices are enrolled by adding them in the dashboard and pointing their DNS-over-HTTPS resolver at the generated address, after which activity appears in the DNS Query Log and Insights.

Enroll a DoH-capable router or gateway into DNSFilter policy without installing any client software — useful for network devices where agent deployment is not possible.

> 📍 In the dashboard, go to Deployments › Clientless and add a new device to receive its unique DoH address. Point the device's DNS-over-HTTPS resolver at that address. Activity will appear in the DNS Query Log and Insights.

— Names exact UI path, plan tier, and setup steps

11/24 November 24 – Web Application and…

02 Unified Audit Logs page Not checked yet NEW

80

A new `Tools > Audit Logs` page replaces the old Policy Audit Log and shows who changed what and when across policies, sites, Roaming Clients (including bulk edits), organization settings, team access, and billing. It can be filtered by audit type, date range, change author, and event, or searched in full, and MSPs can review the change history of any client organization. The old Policy Audit Log address now redirects to this page with Policy pre-selected, so direct links to the old standalone page no longer work.

Review who changed a client org's policy settings and when — useful for MSP change-control audits after an unexpected policy modification.

> 📍 Go to Tools › Audit Logs, filter by audit type Policy, select the relevant date

range, and choose the client organization to review its full change history.

— Exact path, filter options, and migration/redirect behavior given

11/24 November 24 – Web Application and…

03 **Column filters for unblock request queues** Not checked yet

IMPROVED

60

Active Requests and Request History for unblock queues now support column filters by organization, policy, category, block reason, decision, date, and more; exports reflect the applied filters.

— Lists concrete filter fields but no exact UI steps given

11/24 November 24 – Web Application and…

THINNER COVERAGE BELOW

04 **New DNSFilter Root CA G1 certificate** Not checked yet NEW

40

The `Tools > SSL Certificate` download page now adds the DNSFilter Root CA G1 alongside the existing certificate.

— Names exact page and new CA but no further detail

11/24 November 24 – Web Application and…

05 **Roaming Clients API exposes CyberSight status** Not checked yet

IMPROVED

35

The Roaming Clients API now includes CyberSight's upstream connection status.

— Names the API and field but no endpoint or schema detail

11/24 November 24 – Web Application and…

└──▷ **BREAKING ON UPGRADE**

! The old Policy Audit Log address now redirects to `Tools > Audit Logs` with Policy pre-selected; direct links to the old page will no longer open the standalone Policy Audit Log.

WAS THIS USEFUL? 👍 👎

Zscaler

Feedback ⓘ changelog · Nov 24, 2026 BLOG ↗

Zscaler is a cloud-based security platform that provides secure internet access, zero trust networking, and threat protection for users and applications.

Zscaler GovCloud adds MDM-based remote config and Japanese language support for ChromeOS Client Connector 5.0.0.344, major iOS Client Connector 5.0 capabilities including SSID trusted networks and IPv6, and a new Deception permission for VM credential access control.

Zscaler Client Connector 5.0.0.344 for ChromeOS now supports remotely configuring Client Connector features, including administrator-set overrides, through mobile device management (MDM). Not checked yet

Zscaler Client Connector 5.0.0.344 for ChromeOS adds Japanese language support. Not checked yet

Zscaler Client Connector 5.0 for iOS adds Wi-Fi Name (SSID) as a trusted network criterion, TLS 1.3 support with automatic fallback to TLS 1.2, Disaster Recovery for Private Access and Internet & SaaS for Authentication Service, native end-to-end IPv6 from IPv6-only iOS devices to Internet & SaaS and Private Access, and automatic Private Access certificate renewal before expiration. Not checked yet

Zscaler Deception adds a new `Read:VirtualMachines:Credentials` permission to restrict super admins, admins, and users with custom roles from accessing VM credentials and registration tokens. Not checked yet

NEW DETECTION New ThreatParse Rules for Citrix NetScaler ADC and Gateway: CVE-2026-88771

NEW SUPPORTED PLATFORM Real User Monitoring on FedRAMP High in ZDX

WAS THIS USEFUL? 👍 👎

Zscaler Internet Access

Feedback ⓘ

changelog · Nov 24, 2026 BLOG ↗ RANK

Zscaler Internet Access is a cloud-based secure web gateway that inspects and controls internet traffic to protect users from threats and enforce policies.

> Zscaler Internet Access added standardized test FQDNs for validating DNS Control policy behavior.

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

01 **Standardized test FQDNs for DNS Control validation**

Not checked yet **NEW** 40

A set of standardized, immutable test FQDNs with persistent categorization lets admins validate DNS Control policy behavior without relying on external domains that may change over time.

— Describes purpose and mechanism but no exact FQDNs or UI path 2026-11-24-15f1eae2

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

iDenfy

Feedback ⓘ updated · Nov 25, 2026 NOTES ↗ RANK

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

> iDenfy's latest release adds Financial Insights to bank verification, document auto-capture to its mobile SDKs, and expands KYB coverage to three new countries, alongside smaller improvements to blocklisting, AML monitoring and KYB branding.

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

01 **Financial Insights for bank verification** Not checked yet NEW 45

Bank verification now returns Financial Insights — a summary of income, cards and spending — intended for affordability checks.

– Names the data fields but no API/endpoint detail snapshot-20261125

02 **Delete bank verifications** Not checked yet NEW 40

Bank verifications can now be deleted from the API or dashboard.

– Clear action but no endpoint name given snapshot-20261125

03 **Document auto-capture in mobile SDKs** Not checked yet NEW 40

iOS and Android SDKs now support document auto-capture.

– Names platforms but no flag or config detail snapshot-20261125

04 **KYB coverage for Romania, Russia and Latvia** Not checked yet NEW 40

KYB now covers Romanian, Russian and Latvian entities.

– Names three new countries, no mechanism given snapshot-20261125

05 **Blocklist entries without date of birth** Not checked yet IMPROVED 30

You can now add blocklist entries without requiring a date of birth.

– Minimal detail, single requirement removed snapshot-20261125

06 **Custom support contacts on KYB form** Not checked yet NEW 30

You can now show your own support contacts on the KYB form.

— Simple UI customization, no config key named `snapshot-20261125`

07 **Sortable AML monitoring results** Not checked yet IMPROVED

30

AML monitoring results in the dashboard can now be sorted by column.

— Clear UI improvement but minimal detail `snapshot-20261125`

WAS THIS USEFUL? 👍 👎

Didit

Feedback ⓘ updated · Nov 25, 2026 NOTES ↗ RANK

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit streamlined its KYB key-people KYC flow by pre-populating sessions with Document AI data and automatically propagating collected date of birth to child sessions.

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

01 **KYB key-people KYC session prefill and data propagation**

Not checked yet IMPROVED

45

KYB key-people KYC sessions can now be pre-populated from Document AI data, reducing manual re-entry for identity verification flows. Additionally, a date of birth collected during a KYB flow now propagates automatically to each key person's child KYC session.

— Describes behaviour but no config, flag, or endpoint to act on. `snapshot-20261125`

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ announced · Nov 24, 2026 BLOG ↗ RANK

Secure users, AI agents, and more with Auth0, a production-ready authentication and authorization platform built to help you ship faster and drive sign-ups.

> Auth0 now lets Self-Service Enterprise customers purchase additional M2M Token and Enterprise Connection capacity on demand, without a support ticket.

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

01 **Self-service capacity add-ons for M2M Tokens and Enterprise Connections**

Not checked yet NEW 70

Existing Self-Service Enterprise and Professional plan customers upgrading to Enterprise can purchase additional M2M Token capacity blocks and Enterprise Connection capacity blocks (SAML, OIDC) on demand via `Settings > Subscription` in the Auth0 dashboard, scaling backend authentication and federated SSO alongside customer growth.

Add M2M Token or Enterprise Connection capacity blocks without opening a support ticket — useful when backend traffic spikes or a new enterprise customer requires SSO.

> 📍 Log into the Auth0 dashboard › Settings › Subscription, then select the M2M Token or Enterprise Connection add-on blocks that match your required capacity and complete checkout.

— Names exact UI path and plan requirements; no pricing or limit details given

launch-20261124-2f151ced

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ Access profile reviews (2026-11-24)

BLOG ↗ RANK

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

> ConductorOne introduces access profile reviews, a new campaign scope type that lets reviewers audit entitlements, membership rules, and manual enrollments on a per-profile basis, with automated remediation for denied items.

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

01 Access profile reviews campaign scope

Not checked yet NEW

91

Campaigns gain a new By access profiles review scope type, allowing selection of up to 500 profiles per campaign. Reviewers can choose what to decide per profile: profile composition (entitlements granted), membership criteria (auto-enroll rules), or membership exceptions (manual overrides), and can navigate via the new View by access profile option, which shows a tab per review area and lets reviewers reassign with a reason. Denied reviews trigger automated remediation: denied entitlements or membership rules create a change task for the profile owner, denied manual enrollments create a revoke task, and denied membership exclusions automatically lift the exclusion.

Start an access profile review campaign to confirm that high-risk shared profiles still grant only the entitlements they should.

> 📍 In the ConductorOne console, go to Campaigns › New Campaign, set the scope type to By access profiles, select up to 500 profiles, choose Profile composition as the review area, then save as a campaign or template.

— Names exact scope type, limits, review areas, and remediation behavior.

Access profile reviews (2026-11-24)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password for Linux 8.12.48 (2026-11-24)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password's headline addition this window is Environments, a new secrets store for projects and stages that plugs into terminals, CI/CD, and AI coding agents, alongside a public-preview Credential Broker for tokenless CI/CD authentication, new developer security policies, an in-browser Events API tester, and a round of 1Password for Linux refinements including Steam Guard support.

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

01 **New developer security policies in admin Policies settings**

Not checked yet NEW 90

Admins gain several new policy toggles: Allow local disk scanning generates a developer secrets discovery report surfacing plaintext SSH private keys and `.env` files across team members' machines; Allow automatic creation of SSH configuration file auto-creates SSH configuration files and enables SSH Bookmarks; Require team members to add the extension during account setup enforces the browser extension, including SSH item autofill for Git commit signing; a biometric unlock policy under Policies > Authentication > Unlock and auto-lock lets developers authenticate 1Password developer tools with Touch ID, Windows Hello, or Linux system authentication; and non-owner groups can now be granted permission to create service accounts via Permissions > Service Account.



Enforce that all team members install the 1Password browser extension and can autofill SSH public keys during onboarding.

📍 Sign in to 1Password.com › Policies › Manage (under Sharing and permissions) › turn on 'Require team members to add the extension during account setup' › under

Autofill access, confirm SSH items is selected.

Surface plaintext developer secrets (SSH keys, .env files) on team members' local disks via the developer secrets discovery report.

📍 Sign in to 1Password.com › Policies › Manage (under Sharing and permissions) › turn on Allow local disk scanning .

Let a non-owner group (e.g. a DevOps group) create service accounts for CI/CD pipelines without granting full admin rights.

📍 Sign in to 1Password.com › Permissions › Service Account › grant the target group permission to create service accounts.

— Names exact policy toggles with navigation steps for each. product docs

02 **In-browser "Try it" panel for Events API** Not checked yet NEW

85

Each Events API endpoint page (Introspect , Audit events , Item usages , and Sign-in attempts) now has a Try it panel for sending live requests without writing code: paste a bearer token, pick a region (1password.com, 1password.ca, 1password.eu, or ent.1password.com), and send.

Verify a new Events Reporting bearer token and confirm which event types it can read before wiring it into a SIEM pipeline.

📍 Steps: 1. Open the Introspect bearer token endpoint page in the 1Password Events API docs. 2. Select Try it . 3. Choose the server that matches your account (1password.com, 1password.ca, 1password.eu, or ent.1password.com). 4. Paste your token in the Authorization field. 5. Select Send and inspect the live response to see which events the token can read.

— Lists every endpoint, regions, and exact test steps. product docs

03 **Environments integration for AI coding agents** Not checked yet

NEW

75

The Environments MCP server connects coding agents (Claude Code , Codex , Cursor , Kiro , or any MCP client) to project secrets by mounting them as a local `.env` file so nothing sits in plaintext on disk or in MCP config. One-step setup plugins bundle this MCP server with the Environments agent skill for Claude Code , and bundle the agent skill plus a `.env` validation hook that runs before the agent executes shell commands for Cursor . The 1Password Developer Tools plugin for Claude Code manages Environments from within it; the Developer Tools plugin for Cursor manages Environments and validates local configuration. New agent hooks more broadly run

automated checks on a 1Password setup inside supported agents and IDEs to confirm required resources are in place before an agent executes commands.

— Names every client and plugin but gives no runnable command. product docs

04 **Environments secrets store** Not checked yet NEW 70

1Password Environments lets teams create an isolated secrets store for a single project or stage (dev, staging, etc.), replacing per-developer `.env` files so secrets are never written to plaintext. Secrets load into local dev via a mounted `.env` file with nothing written to disk, and can be read in terminal and code via the 1Password CLI or SDKs for Go, JavaScript, and Python. Admins can enable it via the new Policies permission `Allow developers to use Environments to manage their project secrets and use them in tools and workflows`.

Set up a new Environment in the 1Password app to isolate secrets for a specific project stage.

> 📍 In the 1Password app, go to Details › Add your secrets, then import an existing .env file or add key-value pairs manually.

— Names SDKs, CLI, and permission with example setup steps. product docs

05 **Steam Guard codes in one-time password fields** Not checked yet

NEW 65

1Password for Linux can now save Steam Guard codes in one-time password fields, displaying Steam's five-character codes automatically from a pasted `otpauth://` URI.

Add a Steam account's Guard code alongside its password so 1Password generates the five-character Steam OTP automatically.

> 📍 In 1Password, open or create a Login item for your Steam account, add a one-time password field, and paste your Steam Guard `otpauth://` URI — 1Password will display the five-character Steam code.

— Concrete example shows exact setup steps. 1Password for Linux 8.12.48 (2026-11-24)

06 **Credential Broker tokenless CI/CD authentication** Not checked yet

NEW 60

Credential Broker PUBLIC PREVIEW authenticates automated workflows without storing a token by verifying access per run via the CI/CD platform's OIDC identity; it currently supports GitHub Actions and custom OIDC workflows, issuing secrets to agents without a stored token. 1Password also added scoped service-account tokens for automated workflows that need access to specific vaults or Environments.

— Explains mechanism and scope but no setup steps given. product docs

07 **1Password for Agents sign-in approval flow** Not checked yet NEW

60

1Password for Agents lets a personal agent sign in to websites on your behalf; 1Password prompts you for approval and fills username, password, and one-time password into the agent's browser without the agent ever seeing the credentials. A Browserbase integration lets agents running in Browserbase sign in to websites through 1Password for Agents.

— Explains the approval mechanism, no setup steps provided. product docs

THINNER COVERAGE BELOW

08 **Biometric sign-in to AI CLIs via shell plugins** Not checked yet

NEW

45

Sign in to AI CLIs (Claude Code, Codex, Gemini, GitHub Copilot, and more) with biometric authentication using shell plugins.

— Names supported CLIs but no setup instructions. product docs

09 **Developer secrets hygiene tooling** Not checked yet NEW

40

1Password for VS Code detects plaintext secrets in VS Code projects and secures them in 1Password. Developer Watchtower in the desktop app audits team machines for SSH key security issues such as plaintext-stored keys or outdated cryptography.

— Names two tools but describes neither's mechanism. product docs

10 **NVIDIA OpenShell sandboxed agent access (beta)** Not checked yet

NEW

40

Sandboxed agents get time-limited access to secrets via NVIDIA OpenShell BETA, where the agent sees only placeholders instead of real secrets.

— Describes the behavior but no configuration steps. product docs

11 **Keyboard shortcuts settings category** Not checked yet IMPROVED

40

Keyboard shortcuts now have their own category in Settings, found under Settings > Keyboard Shortcuts between Browser and Developer.

— Names exact navigation path, no further detail.

1Password for Linux 8.12.48 (2026-11-24)

12 **Profile relocated to account menu** Not checked yet IMPROVED

40

Profile has moved from the sidebar to the account menu, accessible by selecting your account name at the top of the sidebar.

— Names exact new location, nothing more described.

1Password for Linux 8.12.48 (2026-11-24)

13 **Persistent on-screen notifications** Not checked yet IMPROVED 20

Notifications, like the one that appears when you copy a field, now stay on screen until dismissed.

— Brief behavior change with no additional detail.

1Password for Linux 8.12.48 (2026-11-24)

WAS THIS USEFUL? 👍 👎

Clerk Feedback ⓘ

Set up iOS, macOS, and Android apps with clerk init (2026-11-24) BLOG ↗ RANK

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

> Clerk extended its CLI to support native app setup, letting developers scaffold or configure Clerk in iOS, macOS, and Android projects.

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

01 **Native app setup via clerk init** Not checked yet NEW 67

The `clerk init` command now supports iOS, macOS, and Android projects, allowing developers to set up Clerk in existing Xcode and Android projects or scaffold a new native app with Clerk already configured.

— Names exact command and supported platforms, no further detail on options

Set up iOS, macOS, and Android apps wit…

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Hydra, Ory Kratos, Ory Keto, Ory Polis v26.4.21 released NOTES ↗ RANK

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

> Ory patched a security-relevant matching bug in Hydra's audience validation for OAuth2/OIDC tokens.

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

01 **Exact matching for opaque audience URIs in Hydra** Not checked yet

IMPROVED 50

Ory Hydra now compares opaque audience URIs, such as URNs like `urn:example:service-a`, using exact matching rather than partial matching, preventing tokens from being accepted for audiences they do not precisely match.

— Explains the fix and names the example URI format, no migration steps given

Ory Hydra, Ory Kratos, Ory Keto, Ory Po…

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-11-23) BLOG ↗

RANK

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

> JumpCloud's November 2026 release adds compliance-focused Access Certification campaigns and a private preview of Agent Delegated Access for AI agents, alongside new Windows user-level policies, an Apple Restrictions policy, Device Lock PIN escrow, SCIM group name overrides, and capability-based app setup including MCP-only applications.

└──▷ **WHAT SHIPPED · 9 FEATURES** most completely described first ? what's the number?

01 **Access Certification campaigns for SSO and directory reviews**

Not checked yet NEW 88

Adds Access Certification campaigns for periodic access reviews of SSO applications and directory user groups, with multi-stage approval chains involving administrators, managers, resource owners, or designated user groups. Supports automated remediation that immediately revokes access or suspends inactive identities when a review decision is finalized or a deadline expires, and lets admins export campaign snapshots in CSV, JSON, or PDF formats with SHA-256 document hashing and administrator sign-off flows for compliance audits.

Launch an Access Certification review campaign for an SSO application and export the signed, frozen snapshot for a compliance audit.

> 📍 In the Admin Portal, go to Manage Access Certification, create a new campaign targeting an SSO application, complete the review workflow, then export the campaign snapshot in CSV, JSON, or PDF format with administrator sign-off.

— Full mechanism, formats, and hashing named; includes exact export workflow

Release Notes 2026 (2026-11-23)

02 **Agent Delegated Access for AI agents** Not checked yet NEW

83

Agent Delegated Access PRIVATE PREVIEW lets an AI agent authenticate with its own JumpCloud-managed identity while acting on behalf of an end user, using token exchange between a trusted user identity token and a signed agent-on-behalf-of token. It is configurable via allowed audiences and delegation enablement from the agent's new Delegated access tab in the Admin Portal.

— Token exchange mechanism and config surface named, but preview-only

Release Notes 2026 (2026-11-23)

03 **Device Lock enhancements for Apple MDM** Not checked yet IMPROVED

83

Device Lock for Apple MDM now generates and escrows a lock PIN automatically for macOS, retrievable from the Admin Portal under `Highlights > Device Lock > View PIN`. It adds a `macOS Lock Status` column to the Devices list and supports showing a message and phone number on the Lock Screen for Mac, iPhone, and iPad.

Review and retrieve the escrowed PIN for a locked Mac during an incident response — navigate to the device's lock record in the Admin Portal.

> 📍 In the Admin Portal, go to Highlights › Device Lock › View PIN to retrieve the automatically escrowed lock PIN for a locked Mac.

— Exact UI path, new column name, and retrieval steps given

Release Notes 2026 (2026-11-23)

04 **Per-application SCIM group name overrides** Not checked yet NEW

78

Admins can set per-application SCIM group name overrides using Show In App As on an application's User Groups and Agent Groups tab, choosing Save to store locally or Save and sync to push to the application immediately. Show In App As takes priority over Group Export Attribute Mapping when both define a display name for the same group.

Override the SCIM group name sent to a specific SaaS app without renaming the internal JumpCloud group — useful when an app expects a different naming convention.

> 📍 In the Admin Portal, open the SSO application, go to the User Groups or Agent Groups tab, enter the desired name in Show In App As, then click Save and sync to push the override immediately.

— Exact fields, buttons, and precedence rule with step-by-step example

Release Notes 2026 (2026-11-23)

05 **MCP-driven Workflows management and new templates** Not checked yet

NEW 68

Workflows now supports creating, modifying, duplicating, and troubleshooting JumpCloud Workflows from a connected MCP-enabled AI client using natural language through the existing JumpCloud MCP Server connection. The template library adds ready-to-use templates: Auto-Suspend Users on Inactivity Alert, Auto-Block Newly Discovered Generative AI Apps, `Mac Enrollment - Auto-Revoke Admin Privileges`, `Sensitive Software Install Detection - Email Notification`, `User Group Added - Send Member Email`, and `User Activated - Add to Pre-Hire Holding Group If Before Start Date`.

— Templates named explicitly but no concrete command syntax given

Release Notes 2026 (2026-11-23)

---

06 **User-level Windows policy types** Not checked yet NEW 65

Adds a User Level Windows Custom Registry Keys Policy that applies values under `HKEY_CURRENT_USER` when an assigned user signs in, requiring no MDM enrollment or device restart, and a User Level Custom Windows MDM (OMA URI) Policy for configuring custom MDM settings scoped to a user's profile on MDM-enrolled Windows devices.

— Registry path and OMA URI scope named, no walkthrough given

Release Notes 2026 (2026-11-23)

---

07 **Unified Apple Restrictions policy for macOS and iOS/iPadOS**

Not checked yet NEW 65

JumpCloud adds a single Apple Restrictions Policy covering macOS and iOS/iPadOS devices, with controls spanning system settings, privacy, biometrics, iCloud, and Apple Intelligence & Siri, managed under `Device Management > Policy Management`.

— Named control categories and UI path, no step sequence Release Notes 2026 (2026-11-23)

---

08 **Capability-based app setup with MCP-only applications**

Not checked yet NEW 65

Capability-based app setup lets admins select only the capabilities needed — SSO, SCIM, and/or MCP — when adding an application, and add further capabilities later via the Actions menu. Admins can also create MCP-only applications, assign them to user groups, and make the MCP server available through AI Gateway by selecting Manage MCP Servers during app setup, without configuring SSO.

— Names capabilities and setup control, no exact navigation steps

Release Notes 2026 (2026-11-23)

THINNER COVERAGE BELOW

09 **Independent enforcement controls for macOS auto-updates**

Not checked yet IMPROVED 55

The declarative device management (DDM) policy for Automatic macOS Updates now lets admins set enforcement, grace periods, and local enforcement times for major OS upgrades and minor OS updates independently.

— Describes what changed but no config keys or exact navigation

Release Notes 2026 (2026-11-23)

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Immuta

Feedback ⓘ updated · Nov 25, 2026 NOTES ↗ RANK

Use Immuta to provision governed access to data.

> Immuta exposes whether assets accept access requests directly in search results and tightens validation on the request-submission endpoint to reject calls that were previously silently accepted.

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

01 **Stricter validation on access requests** Not checked yet BREAKING

72

`POST /asset/{id}/request` now rejects two previously-accepted cases: requests against assets where `syncPolicy=false` (disabled parent or leaf assets), preventing approvals that would never be provisioned; and masking exception requests submitted with zero columns or with columns that carry only local masking policies. Callers relying on the old permissive behavior will see new errors on upgrade.

— Endpoint and exact rejection conditions named, no migration steps given

snapshot-20261125

02 **Requestable flag on search results** Not checked yet NEW

70

Asset and column objects returned from search now include a `requestable` flag indicating whether the object accepts access requests, letting callers check eligibility before calling `POST /asset/{id}/request`.

Before submitting an access request, check the `requestable` flag in search results to confirm the asset accepts requests and avoid a rejected call.

```
$ curl -s -X GET 'https://<immuta-host>/api/asset?search=<query>'
-H 'Authorization: Bearer <token>' | jq '.results[] | {id: .id,
name: .name, requestable: .requestable}'
```

— Names the flag and shows a runnable check command snapshot-20261125

└──▷ **BREAKING ON UPGRADE**

! `POST /asset/{id}/request` now rejects requests for assets with `syncPolicy=false` (disabled parent or leaf assets); callers that previously submitted requests against those assets will receive errors on upgrade.

! Masking exception requests with zero columns, or with columns that have only local masking policies, are now rejected by `POST /asset/{id}/request` where they were previously accepted silently.

WAS THIS USEFUL? 👍 👎

Cyera

Feedback ⓘ announced · Nov 23, 2026 BLOG ↗

Cyera helps organizations protect data, govern identities, and secure AI agents to reduce risk, stay in control, and adopt AI with confidence.

Cyera adds Claude Artifacts scanning to Agent Guardian for Claude Enterprise, raising policy issues for sensitive data shared publicly or externally and linking owners to review in Claude.

NEW INTEGRATION NEW SUPPORTED PLATFORM Cyera's Agent Guardian and Data Security coverage now extends to Claude Artifacts in Claude Enterprise, classifying sensitive data in Docs, Slides, Dashboards, Designs, Motion, and code artifacts (apps and websites), and surfacing each artifact's sharing level (owner-only, group, org-wide, or public link). Not checked yet

Cyera policies raise an issue when a sensitive artifact is public (shared with anyone who has the link) or external (shared with someone outside the Claude organization), and each issue includes an Open in Claude link so the artifact owner can review and adjust sharing directly. Not checked yet

Each scanned artifact is added as a data source in the Agent Graph, alongside other AI activity Agent Guardian already maps, giving security teams a unified view of sensitive data across Claude usage. Not checked yet

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

cloudformation-guard

Feedback ⓘ 3.4.0 NOTES ↗ RANK

CloudFormation Guard is an open-source CLI tool that validates AWS CloudFormation templates against policy-as-code rules written in a domain-specific language.

> CloudFormation Guard 3.4.0 fixed YAML boolean parsing and added an NPM distribution channel.

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

01 **NPM package distribution** Not checked yet NEW 40

cloudformation-guard adds an NPM package distribution, making it installable via npm in addition to existing installation methods.

— Names a concrete new install method but no package name given. 3.4.0

02 **Correct parsing of capitalized YAML booleans** Not checked yet

IMPROVED 30

cloudformation-guard now correctly parses capitalized and uppercase YAML booleans (e.g. `True`, `TRUE`) as boolean values instead of mishandling them.

— States the fix but no before/after detail or examples. 3.4.0

WAS THIS USEFUL? 👍 👎

Compyl

Feedback ⓘ Nov 23rd, 2026 NOTES ↗ RANK

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

> Compyl expanded its AWS sync engine with five new resource types — RDS snapshots, Data Lifecycle Manager, Backup, CloudWatch alarms, and Config conformance packs — and added a view-only lock for extension form fields.

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

01 **View-only setting for extension form fields** Not checked yet NEW

65

Extension form fields now support a per-field View only setting that keeps the value visible but prevents editing by anyone, including admins, and limits imports to setting the value only on new record creation.

— Names exact field setting and its scope; UI path implied. Nov 23rd, 2026

02 **CloudWatch alarm sync with history backfill** Not checked yet NEW

60

AWS syncs now collect CloudWatch alarm definitions and alarm history, with a 30-day backfill on first sync and alarm history retained after an alarm is deleted.

— Clear mechanism and number, but no manual action point. Nov 23rd, 2026

THINNER COVERAGE BELOW

03 **AWS sync coverage for RDS, DLM and Config conformance packs**

Not checked yet NEW

57

AWS syncs now collect AWS RDS database snapshots and cluster snapshots, AWS Data Lifecycle Manager policies along with the EBS snapshots those policies create and Auto Scaling groups, and account-level AWS Config conformance packs.

— Names several resource types but no UI path or config key given. Nov 23rd, 2026

04 **AWS Backup protected resources sync** Not checked yet NEW

53

AWS syncs now collect AWS Backup protected resources, refreshed per region so resources no longer backed up are automatically removed.

— Mechanism explained but no manual trigger or setting named. Nov 23rd, 2026

WAS THIS USEFUL?

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ RegScale 6.36.1.1 (2026-11-18) BLOG ↗ RANK

RegScale is a continuous compliance automation platform that helps organizations manage security controls, audit evidence, and regulatory requirements in real time.

> RegScale's only update this window strengthens the Questionnaire module with recurring assessments, record updates from answers, and large-scale results analysis.

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

01 **Questionnaire module recurring assessments and analysis**

Not checked yet IMPROVED 33

The Questionnaire module now supports recurring assessments, can update records directly from questionnaire answers, and allows analyzing results at scale.

— Describes capabilities but no concrete config, flag, or workflow steps

RegScale 6.36.1.1 (2026-11-18)

WAS THIS USEFUL?

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 6.6.0 (2026-11-24) BLOG ↗

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

Runlayer adds MCP run-management tools, subagent scoping, agent reminders, saved browser logins, and Slack approval status cards alongside several hardened credential-handling defaults.

Manage agent runs from MCP clients with new `cancel_agent_run`, `retry_agent_run`, and `rename_agent_run` tools in Runlayer MCP. Not checked yet

List every connector an admin controls from Runlayer MCP using the `list_servers` tool with `scope=manage`. Not checked yet

Let a running agent inspect its own allowed-agent list with the Runlayer MCP `list_agent_allowed_agents` tool. Not checked yet

Share agents, skills, plugins, and artifacts from MCP clients via updated share tools, including `update_skill_shares` which now accepts groups and roles. Not checked yet

Read the tool-rail inset dimensions in HTML artifact apps via CSS variables `--rl-host-chrome-bottom` and `--rl-host-chrome-right` to keep fixed UI clear of the rail. Not checked yet

+13 more

UPDATED DETECTION Private key redaction in sensitive-data and PII detection excerpts

└──▷ BREAKING ON UPGRADE

- ! Slack channel monitors no longer start runs from messages sent by people, bots, or apps outside the organization; enable 'Respond to messages from outside your organization' in the monitor's settings to restore prior behavior.
- ! Runlayer Assistant artifacts are now private by default; artifacts previously open to the whole organization are now restricted — share them from the Share menu to restore access for colleagues.
- ! Microsoft Teams connections now require an admin to enter the organization's Microsoft tenant ID or domain in Settings -> Integrations before connecting; existing connections keep their tenant automatically.

WAS THIS USEFUL? 👍 👎

◆ Agent Sandboxes

NVIDIA OpenShell

Feedback ⓘ v0.1.10 NOTES ↗ RANK

OpenShell runs autonomous AI agents in policy-controlled sandboxes with isolated filesystem, network, and credential access.

> OpenShell v0.1.10 adds offline policy linting and experimental Cedar policy support, a cluster of new credential integrations (workspace secrets, SPIFFE token exchange, GitHub App tokens, Azure workload identity), OTLP metrics export, and richer audit capture, alongside several breaking changes to gateway routing, application URLs, and package/image naming.

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

01 **New credential integrations: workspace secrets, SPIFFE, GitHub App, Azure identity**

Not checked yet NEW 85

Store a shared secret once with `openshell secret create` and reference it across providers via `--material-secret` (enabled via `enable_workspace_secrets` on the gateway); enable SPIFFE-based provider token exchange for shared-supervisor Kubernetes sandboxes via `[openshell.gateway.spiffe_broker]`; mint and rotate GitHub App installation tokens scoped to a repository and permission set without exposing the app private key; and let AKS deployments use Entra workload identity for external Azure PostgreSQL instead of storing a database password.

— Four named config surfaces with clear mechanism each v0.1.10

02 **Offline policy linting and Cedar policy authoring** Not checked yet

NEW 80

Lint policy files offline before deploying them with `openshell policy lint policy.yaml`, and enable Cedar policy authoring and enforcement in development builds with `--features experimental` (existing YAML policies continue to work).

Check a policy file for errors before pushing it to a live gateway, catching mistakes without spinning up a sandbox.

```
$ openshell policy lint policy.yaml
```

— Runnable command and flag given, mechanism explained v0.1.10

03 **Application URLs moved to supervisor origin** Not checked yet

80

BREAKING

Application URLs now come from the supervisor application origin; if you use the app listener, replace `domain` with `base_url`, deploy matching gateway, driver, and supervisor images, recreate existing dedicated sandbox pods with the new supervisor image, and upgrade the CLI.

— Named fields and full migration steps provided v0.1.10

---

04 **Gateway-mediated sandbox routing disabled by default**

Not checked yet

75

BREAKING

Kubernetes Helm deployments no longer serve gateway-mediated sandbox service URLs or HTTP/WebSocket forwarding by default; set `gatewayConfig.'openshell.gateway'.enable_gateway_service_routing: true` before upgrading if you rely on that path.

— Exact Helm key and clear pre-upgrade action given v0.1.10

---

05 **Richer gateway audit records** Not checked yet IMPROVED

70

Capture setting values or exec arguments in gateway audit records by setting `capture_setting_values = true` or `capture_exec_command = true` under `[openshell.gateway.audit]`; audit records now also name every resource touched and flag no-op and replayed changes.

— Named config keys and behavior change, no full workflow v0.1.10

---

06 **Helm gateway and supervisor scaling controls** Not checked yet

NEW

70

Set `gateway.listenAddress` and `supervisor.listenAddress` to `::` to run on IPv6 clusters, and cap shared-supervisor capacity with `supervisor.pool.maxInstancesPerMember` and `supervisor.pool.maxMembersPerGateway` (creates beyond the cap fail with sandbox capacity reached).

— Named Helm keys with explicit failure behavior v0.1.10

---

07 **Packaging rebrand: images, Helm chart, and SDK renamed**

Not checked yet

65

BREAKING

OpenShell images now publish to `quay.io/ai-openshell`; if you set `global.image.registry` to a mirror, the chart now pulls `<mirror>/ai-openshell/<name>` instead of `<mirror>/openshell/<name>`. The Helm chart is renamed from `helm-chart` to `openshell` and now publishes to `oci://quay.io/ai-openshell/charts/openshell` (previously

`oci://ghcr.io/nvidia/openshell/helm-chart`). The TypeScript SDK has moved to `@ai-openshell/sdk` on npmjs.org; `@nvidia/openshell-sdk@0.1.10` is the last release of the old package.

— Every renamed path and package listed, migration implied not spelled out v0.1.10

---

08 **OTLP metrics export from gateway** Not checked yet NEW

60

Enable OTLP metrics export alongside existing Prometheus scraping by setting `enabled = true` under `[openshell.gateway.otlp.metrics]`.

Send gateway metrics to an OTLP collector while keeping existing Prometheus scraping active, so you can view traces and metrics in the same observability backend.

```toml
[openshell.gateway.otlp.metrics]
enabled = true
```

— Exact config key and example given, thin on mechanism v0.1.10

THINNER COVERAGE BELOW

09 **Wildcard credential coverage for API paths** Not checked yet NEW

55

Cover every method under an API path with a single credential by ending a `path_template` value in `/**` (for example `/bot{credential}/**`).

— Named syntax with concrete example, limited mechanism detail v0.1.10

---

10 **OAuth refresh provider field support** Not checked yet IMPROVED

45

OAuth refresh providers now accept `token_endpoint_auth_method=client_secret_basic` and `resource=<uri>` fields.

— Named fields but no context on configuration location v0.1.10

---

11 **Docker gateway sandbox storage and image updates** Not checked yet IMPROVED

40

Docker gateways can now give new workspaces persistent `/sandbox` storage and update a stopped sandbox's image without changing its sandbox ID.

— Describes behavior but no config key or command v0.1.10

---

12 **Sandbox observability improvements** Not checked yet IMPROVED

35

Network and HTTP events now show the actual process behind each sandbox connection instead of pid 0, and startup traces now show the time spent connecting to the boundary as its own span.

— Observed behavior change only, nothing to act on `v0.1.10`

---

13 **Kubernetes 1.33 minimum version requirement** `Not checked yet`

30

`BREAKING`

Kubernetes 1.33 is now the minimum tested version; older clusters are no longer tested and may not work.

— States new minimum version with no migration guidance `v0.1.10`

└──▷ BREAKING ON UPGRADE

! Kubernetes Helm deployments no longer serve gateway-mediated sandbox service URLs or HTTP/WebSocket forwarding by default; set `gatewayConfig.'openshell.gateway'.enable_gateway_service_routing: true` before upgrading if you rely on that path.

! Application URLs now come from the supervisor application origin; if you use the app listener, replace `domain` with `base_url`, deploy matching gateway, driver, and supervisor images, recreate existing dedicated sandbox pods with the new supervisor image, and upgrade the CLI.

! OpenShell images now publish to `quay.io/ai-openshell`; if you set `global.image.registry` to a mirror, the chart now pulls `<mirror>/ai-openshell/<name>` instead of `<mirror>/openshell/<name>`.

! The Helm chart is renamed from `helm-chart` to `openshell` and now publishes to `oci://quay.io/ai-openshell/charts/openshell` (previously `oci://ghcr.io/nvidia/openshell/helm-chart`).

! The TypeScript SDK has moved to `@ai-openshell/sdk` on npmjs.org; `@nvidia/openshell-sdk@0.1.10` is the last release of the old package.

! Kubernetes 1.33 is now the minimum tested version; older clusters are no longer tested and may not work.

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ v0.54.1 NOTES ↗ RANK

> Docker Sandboxes v0.54.1 improves macOS secret cleanup so pre-upgrade secrets can be removed without touching the Keychain manually.

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

01 **macOS secret removal without manual Keychain cleanup**

Not checked yet **IMPROVED** 35

On macOS, secrets saved before an upgrade can now be removed without manually deleting their Keychain entries.

— Describes behavior but no command or UI path given v0.54.1

WAS THIS USEFUL? 👍 👎

microsandbox

Feedback ⓘ v0.8.2 NOTES ↗ RANK

> microsandbox shipped a hosted cloud backend that runs existing code remotely via two environment variables, a new secrets API that keeps API keys out of the sandbox guest, plus image-pruning exclusions, several SDK-level sandbox isolation and staging options, and a browser-based root shell in the dashboard.

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

01 Secrets API with `.secret()` builder

Not checked yet NEW

91

A new secrets API uses the .secret() builder method with .env(), .value(), and .allow() to keep API keys out of the sandbox guest — the real value is swapped in only on requests to allowed hosts, and is stored as a placeholder such as `$MSB_GITHUB_TOKEN` inside the VM.

Keep a secret API key out of the sandbox guest and restrict its use to a single allowed host.

javascript

```
import { Sandbox } from "microsandbox";

const sb = await Sandbox.builder('secrets-demo')
  .image('python:3.12')
  .secret((s) =>
    s.env('GITHUB_TOKEN')
     .value(process.env.GITHUB_TOKEN!)
     .allow('api.github.com')
  )
  .create();
```

— Named builder methods and placeholder mechanism with code example.

launch-20261124-c5145c02

02 Hosted cloud backend via environment variables

Not checked yet NEW

85

Existing microsandbox code can now run on hosted cloud infrastructure by setting `MSB_BACKEND=remote` and `MSB_API_KEY` — no code changes, no separate daemon or client required.

Point an existing microsandbox script at the hosted cloud backend without changing any application code.

```
$ export MSB_BACKEND=remote
export MSB_API_KEY="msb_..."
npx tsx job.ts
```

— Exact env vars and a runnable command given. launch-20261124-c5145c02

03 **`--exclude-from` flag for image pruning** Not checked yet NEW

75

`msb image prune` now accepts `--exclude-from`, letting you protect a named list of images from a bulk prune, e.g. keeping CI base images intact between job runs.

Keep specific images safe from a bulk prune — useful in CI where base images must survive between job runs.

```
$ msb image prune --exclude-from ./images-to-keep.txt
```

— Exact flag and subcommand with runnable example. v0.8.2

THINNER COVERAGE BELOW

04 **Metric snapshots read directly from registry files**

Not checked yet NEW

45

Independent metric snapshots can now be read directly from registry files without requiring a running metrics flush cycle, making metrics inspectable even when the flush cycle isn't active.

— Explains mechanism but names no command or endpoint. v0.8.2

05 **Browser-based root shell in dashboard** Not checked yet NEW

45

The dashboard provides a browser-based root shell accessible seconds after creating a sandbox, with no SSH keys required.

— UI path given but no further mechanism detail. launch-20261124-c5145c02

06 **Independent file materialization for snapshots** Not checked yet

NEW

30

Snapshots can opt in to independent file materialization, letting each snapshot resolve its own file copies rather than sharing them.

— Describes behaviour but names no flag or config key. v0.8.2

07 **New SDK configuration capabilities** Not checked yet NEW

30

The SDK adds an opt-in strict mode for host file opens (controlling how the host resolves files for sandboxes), new runtime isolation settings for configuring sandboxes with additional isolation options directly from code, and bounded patch source staging for staging patch sources within bounded limits.

— Prose description only, no named setting. No concrete option names or mechanism given. No concrete limit values or flag named.

v0.8.2

08 **Explicit in-band relay connections for agents** Not checked yet

NEW

20

Agents can now make explicit in-band relay connections, controllable through the agent configuration.

— No specific config keys or mechanism named. v0.8.2

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Security Posture & Discovery

Zenity

Feedback ⓘ announced · Nov 25, 2026 BLOG ↗ RANK

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

> Zenity launched its MCP Gateway as generally available, enforcing identity, inline policy, and kill-switch control over every Model Context Protocol tool call, alongside session-aware threat detection and shadow MCP discovery.

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

01 **Govern MCP server access by name or URL** Not checked yet NEW

70

Govern MCP server access by server name or upstream target URL, allowing approved servers to run freely while blocking or restricting unapproved ones to read-only tools.

Restrict contractor access to a shared Atlassian MCP server so they can read but never write — without deploying a separate server or changing agent config.

> 📍 Steps in Zenity console: 1. Go to MCP Gateway › Servers and select your Atlassian connector. 2. Under Boundaries, add a rule: Identity = contractor (role or domain), Tool agency = mutative or destructive, Action = Block. 3. Add a second rule: Identity = employee, Tool agency = any, Action = Allow. 4. Save — both rules take effect on the next tool call, no endpoint changes required.

— Includes exact console steps to restrict a server.

Zenity MCP Gateway Is Generally Availab…

02 **Inline policy enforcement via Boundaries** Not checked yet NEW

65

Inline policy (Boundaries) evaluates every tool call before it executes, combining the caller's identity (domain, role, permissions) with the tool's agency (read-only, mutative, or destructive) to block disallowed actions. The same Boundaries used for coding agents, Claude Enterprise, and Copilot Studio apply to MCP Gateway traffic, so policy is written once and enforced everywhere.

— Explains mechanism and scope but no setup example.

Zenity MCP Gateway Is Generally Availab…

03 **Session-aware Boundaries for multi-step attacks** Not checked yet

NEW

60

Session-aware Boundaries track earlier calls within a session so risky multi-step sequences — such as pulling financial data then sending a message — are blocked the moment they form, even when each individual call looks legitimate, stopping a workflow the instant it turns into an exfiltration path.

— Clear mechanism and example scenario, no config steps. launch-20261125-d9cacb9b

Zenity MCP Gateway Is Generally Availab…

04 **Kill switch to cut off MCP servers** Not checked yet NEW

60

A kill switch cuts off any MCP server for every client and every user within seconds, without touching endpoint configuration.

Immediately revoke access to a compromised MCP server vendor across every agent and user without touching individual client configs.

📍 Steps in Zenity console: 1. Go to MCP Gateway › Servers. 2. Find the affected server by name or upstream target URL. 3. Toggle the kill switch to Off. 4. Open Activity to review all recorded tool calls (user, AI client, tool, arguments) made before the cutoff.

— Has exact console steps to revoke access instantly.

Zenity MCP Gateway Is Generally Availab…

THINNER COVERAGE BELOW

05 **MCP Gateway for every MCP tool call** Not checked yet NEW

45

Zenity's MCP Gateway sits between AI agents and MCP servers, enforcing identity, inline policy, and kill-switch control over every MCP tool call. It needs no endpoint installation and works with any remote-MCP client.

— Describes capability but no setup steps given. Zenity MCP Gateway Is Generally Availab…

launch-20261125-d9cacb9b

06 **Identity attribution for MCP calls** Not checked yet NEW

45

Ties every MCP tool call to a real user identity through your identity provider, so shared API keys no longer obscure who made a call.

— Names the mechanism but no steps to configure it.

Zenity MCP Gateway Is Generally Availab…

07 **Shadow MCP server discovery** Not checked yet NEW 45

Shadow MCP detection identifies MCP servers configured across the environment — including agent calls that bypass the gateway — before they become a governance gap,

letting teams triage them from a single console.

— Describes detection but no walkthrough of triage console. launch-20261125-d9cacb9b Zenity MCP Gateway Is Generally Availab…

08 **Central activity log for MCP calls** Not checked yet NEW

40

Records every tool call with the user, AI client, tool, and arguments in a central activity log alongside the rest of your agent activity.

— Names log fields but no retrieval steps given. Zenity MCP Gateway Is Generally Availab…

WAS THIS USEFUL? 👍 👎

Harmonic Security

Feedback ⓘ announced · Nov 25, 2026

BLOG ↗

Harmonic Security detects shadow AI and prevents sensitive-data leaks across employee web, desktop, and agent workflows.

Harmonic Security adds a CrowdStrike Falcon Next-Gen SIEM integration to surface AI-driven data exposure and unsafe agent actions in the SOC.

NEW INTEGRATION CrowdStrike Falcon Next-Gen SIEM integration

WAS THIS USEFUL? 👍 👎

◆ AI Red Teaming & Testing

clawscan

Feedback ⓘ v0.4.0 NOTES ↗ RANK

> clawscan v0.4.0 adds an aig-mcp plugin scanner, an optional Decisions API judge, and a pluggable post-verdict overview step, while overhauling judge behavior and workspace layout across several breaking changes.

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

01 **Decisions API judge command** Not checked yet NEW 95

Adds the optional `clawscan-decisions-judge` command to the runtime image, selectable with `--judge`, which calls an OpenAI Decisions API judge (`gpt-6-luna`) and returns a typed `verdict`, `confidence`, `dimensions.<name>.status` / `.confidence`, `scan_findings_in_context[]`, and a `decision` object with model, usage, latency, verdict probabilities, and risk-signal probabilities.

Override the default evidence-only judge with the Decisions API judge to get a full `decision` object including model, usage, latency, and risk-signal probabilities.

```
$ clawscan --profile clawhub --judge clawscan-decisions-judge
<target>
```

— Full response schema and invocation flag given. v0.4.0

02 **Pluggable overview.command post-verdict step** Not checked yet

NEW

95

Adds an optional pluggable `overview.command` profile setting, configurable via `--overview <cmd>` or `--no-overview`, that runs after a completed verdict with `judge.json` and the `{{ judge.result }}` prompt placeholder and records `overview` in the run artifact without affecting verdict, gate, status, or exit status. The `clawscan-decisions-judge` command now moves `summary`, dimension `detail`, finding `note`, and `user_guidance` out of the judge result into this new top-level `overview.result`.

Attach a custom overview command that runs after a verdict and records prose in the run artifact without touching the verdict or exit status.

```
$ clawscan --profile clawhub --overview my-overview-cmd <target>
```

— Exact flags, placeholder and field relocation specified. v0.4.0

03 **aig-mcp plugin scanner for OpenClaw targets** Not checked yet NEW

90

Adds the `aig-mcp` plugin-only scanner (Tencent `A.I.G` `aig-mcp-scan` 0.2.0) and the `pluginScanners` profile field; the `clawhub` profile now routes OpenClaw plugin targets to `aig-mcp` instead of `scanners.skillspector` / `scanners.aig` (now `scanners["aig-mcp"]`). Host runs require the `AIG_MCP_SCAN_ALLOW_HOST_EXEC=1` environment variable, and plugin-target scanner commands now get a 30-minute timeout budget, up from 20 minutes.

Run a plugin scan on an OpenClaw plugin target on a host machine, enabling the new `aig-mcp` scanner with the required environment variable.

```
$ AIG_MCP_SCAN_ALLOW_HOST_EXEC=1 clawscan --profile clawhub
<plugin-target>
```

— Names scanner, env var, config field and timeout change. v0.4.0

04 **Evidence-only judging for ClawHub profiles** Not checked yet

85

BREAKING

The `clawhub` and `clawhub-aig` profiles now judge with an evidence-only Codex run (`gpt-6.1-sol`) controlled by the new `judge.evidenceOnly` / `overview.evidenceOnly` profile setting, adding prompt placeholders `{{ scanner_results }}` and `{{ run_context }}`. This removes the artifact copy

from the judge workspace, refuses the `{{ target.files }}` placeholder, and means `judge.result` no longer carries the Decisions `decision` record.

— Config keys and removed placeholder/field named precisely. v0.4.0

05 **Judge-agnostic runner removes artifact-inspection receipt**

Not checked yet DEPRECATED 75

The runner is now judge-agnostic: it removes the artifact-inspection receipt system entirely, including the challenge file, receipt validation, the `{{ artifact_inspection.required_file }}` placeholder, the `artifact_inspection` field, and the `challenge`, `requiredFile`, and `requiredFileSha256` fields on `judge.attempts`. The runner-level judge prompt budget and all profile-name special cases are also removed.

— Lists every removed field but gives no migration command. v0.4.0

06 **Unified judge workspace layout** Not checked yet BREAKING 65

All profiles now share one judge workspace layout, writing `scanners/<id>.json` and `metadata.json`, with an `artifact/` directory present only when `evidenceOnly` is disabled; the ClawHub-only `artifact/` scan copy is gone.

— Exact file paths given but no migration steps. v0.4.0

07 **virustotal and clawscan-static scanners removed** Not checked yet DEPRECATED 60

The `virustotal` and `clawscan-static` scanners are removed from clawscan; consumers that read `scanners["clawscan-static"]` must switch to another eligibility signal.

— Names removed scanners and required migration key. v0.4.0

└──▷ **BREAKING ON UPGRADE**

! The `clawhub` and `clawhub-aig` profiles now use evidence-only judging (`judge.evidenceOnly` / `overview.evidenceOnly`): the artifact copy is removed from the judge workspace, `{{ target.files }}` is refused, and `judge.result` no longer carries the Decisions `decision` record.

! Plugin targets: consumers that read `scanners.skillspector` or `scanners.aig` for plugin runs must switch to `scanners["aig-mcp"]`; host runs additionally require the `AIG_MCP_SCAN_ALLOW_HOST_EXEC=1` environment variable.

! The `clawscan-decisions-judge` command moves `summary`, dimension `detail`, finding `note`, and `user_guidance` out of judge result to the new top-level `overview.result`.

! The runner is now judge-agnostic: the artifact-inspection receipt is removed (challenge file, receipt validation, `{{ artifact_inspection.required_file }}` placeholder, `artifact_inspection` field, and the `challenge`, `requiredFile`, and `requiredFileSha256`

fields on `judge.attempts`); the runner-level judge prompt budget and all profile-name special cases are also removed.

! All profiles now share one judge workspace layout (`scanners/<id>.json` and `metadata.json`; `artifact/` only when `evidenceOnly` is disabled); the ClawHub-only `artifact/` scan copy is gone.

! The `virustotal` and `clawscan-static` scanners are removed; consumers that read `scanners["clawscan-static"]` must switch to another eligibility signal.

WAS THIS USEFUL? 👍 👎

◆ Model & Agent Protection

ToolHive

Feedback ⓘ v0.69.0 NOTES ↗ RANK

> ToolHive's v0.69.0 release adds EdDSA-signed token support and Microsoft Entra on-behalf-of backend authentication, extends observability into vMCP code mode, and tightens its embedded authorization server by rejecting RFC 8693/JWT-bearer requests that carry an `authorization_details` claim.

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

01 Rejection of authorization_details claims in token exchange

Not checked yet

83

BREAKING

ToolHive's embedded authorization server now refuses RFC 8693 token-exchange requests whose `subject_token` or `actor_token` carries an `authorization_details` claim, returning `400 invalid_request`; it also refuses JWT-bearer or ID-JAG assertions carrying the claim, returning `400 invalid_grant` or `400 invalid_assertion` without consuming the assertion. Previously the claim was silently dropped and these requests succeeded, so a trusted identity provider that puts `authorization_details` (RFC 9396) into tokens or ID-JAGs sent to ToolHive will now have those requests fail.

– exact error codes and claim named, but no migration guidance v0.69.0

THINNER COVERAGE BELOW

02 Observability for vMCP code mode

Not checked yet IMPROVED

55

ToolHive's vMCP code mode now emits OTel metrics, spans, and logs, and reports tool calls in anonymous usage metrics, giving operators visibility into code-mode activity.

– describes what changed but no metric names or config v0.69.0

03 Entra on-behalf-of backend authentication

Not checked yet NEW

50

ToolHive adds Microsoft Entra on-behalf-of (OBO) backend authentication in the open-source edition, enabling delegated identity flows to downstream APIs.

– names the feature but no setup steps or flags v0.69.0

04 **EdDSA-signed token support** Not checked yet NEW 40

ToolHive's token validator now accepts `EdDSA` `-signed` tokens, widening the set of identity providers that can authenticate to the gateway.

— brief mention, no config or flag to act on `v0.69.0`

└──▷ **BREAKING ON UPGRADE**

! The embedded authorization server now rejects RFC 8693 token-exchange requests whose `subject_token` or `actor_token` carries an `authorization_details` claim with `400` `invalid_request`; previously the claim was silently dropped and the request succeeded. If a trusted identity provider puts `authorization_details` (RFC 9396) into tokens or ID-JAGs it sends to ToolHive, those requests now fail.

! JWT-bearer or ID-JAG redemption requests whose assertion carries an `authorization_details` claim now return `400 invalid_grant` or `400 invalid_assertion` and the assertion is not consumed; previously they succeeded.

WAS THIS USEFUL?  

tenuo Feedback ⓘ v0.3.8 NOTES ↗ RANK

> tenuo's v0.3.8 release adds Python-side spending controls — host-policy argument ceilings and token budgets — plus Rust AAT JWT/JWS chain verification, TypeScript warrant inspection tools, and expanded delegation narrowing rules.

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

01 Token budget metering for LLM calls Not checked yet NEW

95

Meter model calls against a warrant's token budget using tenuo.budget_http.install() to intercept all `httpx` / `httpx2` clients in a process, or metered_client() for a single client; budgets are expressed as a Range constraint on an `llm.chat` capability via `tenuo.budget.llm_capability` (Python).

Meter every httpx model call in a Python agent process against a token budget carried in the active warrant — no per-call changes needed.

python

```
import tenuo.budget_http
import tenuo.budget
from tenuo import Capability, Range, configure, mint_sync,
SigningKey

configure(issuer_key=SigningKey.generate(), dev_mode=True,
audit_log=False)

# Install the meter once at startup; it intercepts all
httpx/httpx2 clients.
tenuo.budget_http.install()

# Mint a warrant with a 1000-token ceiling on llm.chat calls.
with mint_sync(
    tenuo.budget.llm_capability(Range.max_value(1000)),
    ttl=300,
):
    # Any httpx call to a model endpoint is now authorized and
metered.
    import httpx
    httpx.post('https://api.example.com/v1/chat', json={...})
```

— Named functions, constraint type, and a runnable install snippet. v0.3.8

02 **Per-tool host-policy ceilings in Authorizer** Not checked yet NEW

95

Apply a per-tool argument ceiling at the enforcement point, on top of any warrant, using Authorizer.set_host_policy(tool, allow) in Python; receipts carry the policy digest as `policy_definition_hash` (key 11), matching the TypeScript SDK's verify(..., { allow }) parameter. Host-policy expressions are now normalized using normalized path roots, omitted defaults, and compiled aliases, so `policy_definition_hash` values for newly issued receipts can differ from earlier SDKs for noncanonical expressions (existing signed receipts remain verifiable).

Enforce a per-tool argument ceiling at the authorizer level in Python, independent of whatever the incoming warrant says, so a rogue or over-scoped warrant cannot exceed the host's local policy.

python

```python
from tenuo import Authorizer, Range

authorizer = Authorizer()

# Cap 'scale' at 5 replicas regardless of warrant contents.
authorizer.set_host_policy('scale', allow={'replicas':
Range.max_value(5)})

# Receipts issued through this authorizer will carry
policy_definition_hash (key 11).
```

— Exact method signature, receipt field, and a compatibility caveat given. v0.3.8

03 **AAT verification in JWT/JWS encoding (Rust)** Not checked yet NEW

85

Verify Attenuating Agent Tokens in JWT/JWS encoding in Rust using the new `tenuo::aat` module with the `aat-jws` feature flag; `verify_chain` runs the chain verification algorithm of `draft-niyikiza-oauth-attenuating-agent-tokens-02` Section 8 and passes the draft's 229 test vectors.

Verify an Attenuating Agent Token chain in JWT/JWS encoding in a Rust service without issuing or deriving new tokens.

toml

```toml
# Cargo.toml
[dependencies]
tenuo = { version = "0.3.8", features = ["aat-jws"] }
```

— Precise spec reference and test-vector count, but only a dependency snippet. v0.3.8

04 **Issuer-managed period budgets via budget_lease** Not checked yet

NEW 78

Spend issuer-managed period budgets (day, week, or month) with the experimental `tenuo.budget_lease` module; call install(lease=LeasedBudget(CloudLeases(trigger_id))) to take short-lived warrant slices and report each model call's usage as an authorization event to Tenuo Cloud (Python).

— Named module and call pattern but no runnable example. v0.3.8

---

05 **Warrant inspection and explanation in TypeScript** Not checked yet

NEW 68

Decode and inspect warrants without possession checks in TypeScript using createTenuo.inspectWarrant() to decode unverified contents and tenuo.explainWarrant() to report chain, constraint, and approval evidence; neither produces an authorization verdict or receipt.

— Named functions and behaviour, but no usage example. v0.3.8

THINNER COVERAGE BELOW

06 **Expanded delegation argument-narrowing rules** Not checked yet

48

BREAKING

Delegation now accepts two additional ways to narrow a tool's arguments: a child may constrain an argument the parent admits without declaring it, and a child that admits no undeclared arguments may omit one the parent marks optional. Upgrade verifiers before issuing chains that rely on these rules, since older verifiers reject chains using the new narrowing patterns.

— Behaviour described but no named flags or migration command. v0.3.8

└──▷ **BREAKING ON UPGRADE**

! Upgrade verifiers before issuing chains that use the expanded argument-map narrowing rules; older verifiers reject chains where a child constrains an undeclared parent argument or omits an optional one the parent marks optional.

! Normalized host-policy expressions now use normalized path roots, omitted defaults, and compiled aliases; `policy_definition_hash` values for newly issued receipts can differ from earlier SDKs for noncanonical expressions (existing signed receipts remain verifiable).

WAS THIS USEFUL? 👍 👎

Minimal Software Research minimal

Feedback ⓘ

release-75bea6c1 NOTES ↗ RANK

Minimal is a secure, declarative runtimes for Claude Code, Gemini, Codex, Pi and more, so you can run multiple agents safely on your stack from a verified software supply chain.

> Minimal ships a Box Egress Proxy with audit and network inspection tooling, adds secret management and GitHub sign-in commands, and introduces opt-in OpenTelemetry tracing across the CLI, daemon and VM host — alongside a breaking CLI cleanup that removes two commands.

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

01 **Box Egress Proxy audit and network inspection** Not checked yet

NEW 80

Minimal adds a Box Egress Proxy governing outbound network access from boxes. `min box audit mybox --parent` reviews the full audit trail for a box, merging in its parent chain, to verify no proxy decision has been tampered with, and `min box show mybox --network` inspects the effective networking and egress policy for a running box before trusting it with outbound access.

Review the full audit trail for a box, merging in its parent chain, to verify no proxy decision has been tampered with.

```
$ min box audit mybox --parent
```

Inspect the effective networking and egress policy for a running box before trusting it with outbound access.

```
$ min box show mybox --network
```

— Runnable commands shown but underlying proxy mechanism unexplained release-75bea6c1

02 **OpenTelemetry tracing across CLI, daemon and VM host**

Not checked yet NEW 78

Adds opt-in OpenTelemetry export with a local spool via `mlog`, propagating one trace per request across the CLI and the daemon (`min` and `minimald`), and exports VM host daemon spans, carrying telemetry into the guest, via `minvmd`. `min bug` now carries the telemetry spool and redacts telemetry settings from diagnostics.

— Mechanism and components named but no enabling command given release-75bea6c1

03 **Removed `min stop` and `min task run --path`** Not checked yet

70

BREAKING

The bare `min stop` and `min task run --path` commands are removed; use `min host stop` and `-C <dir>` instead.

— Exact migration path from old to new commands given release-75bea6c1

THINNER COVERAGE BELOW

04 **GitHub sign-in with `min login`** Not checked yet NEW

55

Run `min login` as the GitHub sign-in command; logging out deletes the sign-in and refuses every reference to it.

— Command named, logout behaviour described, mechanism thin release-75bea6c1

└──▷ BREAKING ON UPGRADE

! The bare `min stop` and `min task run --path` commands are removed; use `min host stop` and `-C <dir>` instead.

WAS THIS USEFUL? 👍 👎

◆ AI Runtime Defense & Guardrails

Manifold Security

Feedback ⓘ announced · Nov 19, 2026

BLOG ↗ RANK

Manifold Security is an AI detection and response platform that monitors what autonomous AI agents do at runtime on enterprise endpoints and blocks risky actions.

> Manifold Security's launch adds cost visibility and control for AI coding agents, combining spend anomaly alerts, cross-vendor per-person budgets with warn/block enforcement, MCP access restrictions on overage, and finance-ready spend exports.

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

01 Per-person AI agent budgets with warn/block enforcement

Not checked yet NEW 78

Set one budget per person that spans every coding agent (Claude Code, Codex, Cursor, and others), so a cap set once applies across all vendors. Budgets can be applied to individuals, to a team from Microsoft Entra or custom tags, or to a single vendor, agent, or model. Enforcement follows a warn-at-80% `-then-block-at-100` % policy, delivering a warning inside whichever agent the person is using before stopping the next prompt across all agents.

Walk through setting up a warn-then-block budget policy for a team to catch runaway agent loops before they become a large bill.

> 📍 In the Manifold console, go to Policies › Budgets, create a new budget, set the monthly limit, assign it to a Microsoft Entra team or a custom tag, then set Warn at 80% and Block at 100% as the enforcement responses.

— Mechanism and scope detailed; console steps given but no API/flag names

launch-20261119-32254855

THINNER COVERAGE BELOW

02 Premium MCP server and skill access limits on budget overage

Not checked yet NEW 44

Once a person exceeds their budget, Manifold limits access to premium MCP servers and skills, using the same policy response system as existing Manifold agent rules.

— Mechanism named but no specific rule or config key given launch-20261119-32254855

03 Spend anomaly alerts with full session context Not checked yet

NEW

38

Flag abnormally high agent token usage with spend alerts that include the full session so you can see what the agent was doing while costs climbed.

— Describes trigger and payload but no config surface named `launch-20261119-32254855`

04 **Spend export for finance reporting** Not checked yet NEW

36

Export spend broken down by team, agent, or model, priced at negotiated rates, for finance reporting.

— Export capability named but no format or endpoint specified `launch-20261119-32254855`

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

WAF - Updated prompt logging for AI Security for Apps (2026-11-24) BLOG ↗

RANK

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

> Cloudflare shipped a new payload-logging workflow for AI prompt security in the WAF, alongside a CLI feedback command, AVIF conversion in Polish, encoded range-request forwarding, email-based 2FA verification for social logins, and free email sending on the Workers Free plan.

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

01 **Payload logging for AI prompt content in WAF** Not checked yet

88

BREAKING

Prompt content can now be captured by configuring payload logging and writing a WAF custom rule whose expression references a field with the `cf.llm.prompt` prefix (e.g. `cf.llm.prompt contains "ignore previous instructions"`), configured under Security > WAF > Custom Rules with payload logging enabled on the zone. This replaces the previous AI Security Log Mode Ruleset procedure for logging prompt content.

Capture extracted AI prompt content in payload logs when a custom WAF rule fires on an LLM request.

> 📍 In the Cloudflare dashboard, go to Security › WAF › Custom Rules and create a rule whose expression references a `cf.llm.prompt` field (e.g. `cf.llm.prompt contains "ignore previous instructions"`), then enable payload logging on the same zone to store the matched prompt.

— Names exact field prefix, rule example, and dashboard path

WAF - Updated prompt logging for AI Sec…

02 **CLI feedback command** Not checked yet NEW

70

Users can report bugs and feature requests directly from the Cloudflare CLI using `cf feedback`, or via the `submit_feedback` MCP tool.

Send a bug report or feature request directly from your terminal without leaving the CLI.

```
$ cf feedback
```

– Exact runnable command given product docs

THINNER COVERAGE BELOW

03 **Accept-Encoding forwarding on range requests** Not checked yet

IMPROVED

55

Uncached HTTPS GET range requests now forward the client's Accept-Encoding header, and cached encoded ranges are served only to clients that accept that encoding.

– Clear mechanism but automatic, nothing to configure product docs

04 **AVIF conversion in Polish** Not checked yet NEW

50

Polish now converts JPEG and PNG origin images to AVIF and serves the smaller format to browsers that support it.

– Names formats but no config surface to act on product docs

05 **Email verification for 2FA on social logins** Not checked yet NEW

45

Users who sign in with Apple, GitHub, or Google can now verify their identity by email when enrolling in two-factor authentication, without needing a password.

– Names providers but no exact enrollment steps given product docs

06 **Free email sending on Workers Free plan** Not checked yet NEW

15

Cloudflare adds free email sending on the Workers Free plan.

– Mentioned only in summary with no further detail product docs

└──▷ BREAKING ON UPGRADE

! The previous AI Security Log Mode Ruleset procedure for logging prompt content is replaced by the new payload logging + `cf.llm.prompt` custom rule workflow.

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

Arcade

Feedback ⓘ changelog · Nov 25, 2026 CHANGELOG ↗ RANK

Arcade is the enterprise MCP gateway, built for authorization, integration, run-time policies, and air-gapped deployment.

> Arcade's November 25 release adds multi-account tool profiles, IAM-token PostgreSQL authentication, a full project-data purge, and a new Snipe-IT toolkit, alongside CLI controls for organization suspension and sign-up domain blocking, dashboard search/filtering, API key scoping, and several breaking removals including account-bound gateways and the Porter deployment provider.

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

01 **CLI admin commands: suspension, domain blocking, log purge**

Not checked yet NEW 85

Adds `arcade organizations suspend <org-id>` (and unsuspend) to suspend or unsuspend organizations and accounts and remove a suspended organization's deployed workers in one command, `arcade signup-domains block <domain>` (and unblock) to block or unblock sign-up domains — sign-ups are already refused from 100,000+ disposable email domains and domains without valid mail records — and `arcade executions purge-logs` to purge old tool-execution logs on demand, alongside a new per-project cap on pending deployments.

Suspend a compromised organization immediately to revoke all its credentials and stop its workers from accepting calls.

```
$ arcade organizations suspend <org-id>
```

Block a domain so users can no longer sign up with addresses from it, reducing exposure from disposable or malicious email providers.

```
$ arcade signup-domains block <domain>
```

Purge stale tool-execution logs to reduce Engine storage and limit the blast radius of a data exposure.

```
$ arcade executions purge-logs
```

— Runnable command with clear mechanism and scope. Named command plus concrete numeric threshold. Named CLI command, but cap value unspecified.

changelog-20261125-bdd17c1c

02 **Full project-data purge on deletion** Not checked yet NEW

80

Deleting a project now erases all Engine-held data — end-user accounts, OAuth tokens, secrets, gateways, deployed workers and servers, and tool-execution history — after a configurable retention window; a CLI command purges projects deleted before this release.

— Detailed scope and before/after, command not named exactly. changelog-20261125-bdd17c1c

03 **Automatic schema migrations on helm upgrade** Not checked yet

80

BREAKING

Schema migrations now run automatically on `helm upgrade` unless `<service>.database.migrations` is set to `false`; services that depend on new columns/tables (suspend/unsuspend, key deactivation, domain blocklist, Engine purge job) require the migrations to run.

— Names exact config key and lists dependent features. changelog-20261125-bdd17c1c

04 **Tool execution visibility and usage event reporting**

Not checked yet IMPROVED

68

Tool executions now show the caller's client IP, can be filtered by API key and gateway, and report whether the call ran against a deployed server; usage reporting also sends new events for multi-profile management actions, submitted registry requests, and a deployed server's release finishing.

— Several named filters and event types, no API surface given.

changelog-20261125-bdd17c1c

05 **Multi-account tool profiles for agents** Not checked yet NEW

63

Agents can now work against more than one authorized account for the same tool, called profiles; organizations opt in with a multi-profile switch, and agents can list, name, and switch profiles.

— Behaviour explained but no exact command or flag. changelog-20261125-bdd17c1c

06 **failure_source field in API error responses** Not checked yet

IMPROVED

62

API error responses now carry a `failure_source` field naming who can act on the error — client, end user, organization, third party, or platform.

— Named field and enum values, no endpoint given. changelog-20261125-bdd17c1c

07 **IAM-token PostgreSQL authentication** Not checked yet NEW

62

The Engine, Coordinator, and experience API can now authenticate to PostgreSQL with a short-lived IAM token instead of a stored password.

— Named components, no config key provided. changelog-20261125-bdd17c1c

THINNER COVERAGE BELOW

08 **Project API key deactivation and scoping** Not checked yet IMPROVED

55

Project API keys can be deactivated or reactivated and scoped to calling tools only; new keys default to calling tools.

— Clear behaviour change, no endpoint or flag named. changelog-20261125-bdd17c1c

09 **Dashboard search and filtering for gateways, secrets, OAuth providers**

Not checked yet IMPROVED

53

Gateways, secrets, and OAuth providers now support server-side filter, sort, search, and paging on their dashboard pages.

— Names three affected dashboard areas, no exact path. changelog-20261125-bdd17c1c

10 **Porter deployment provider removed** Not checked yet BREAKING

53

The Porter deployment provider is removed; deployments now default to Kubernetes.

— States removal and new default, no migration detail. changelog-20261125-bdd17c1c

11 **Account-bound gateways removed** Not checked yet BREAKING

48

Support for account-bound gateways is removed; existing account-bound gateway configurations will break on upgrade.

— States the break but no migration steps given. changelog-20261125-bdd17c1c

12 **busybox support image removed** Not checked yet BREAKING

36

The `supportImages.busybox` support image is removed.

— Bare removal notice naming the config key, no replacement. changelog-20261125-bdd17c1c

└──▷ BREAKING ON UPGRADE

- ! Support for account-bound gateways is removed; existing account-bound gateway configurations will break on upgrade.
- ! `supportImages.busybox` support image is removed.
- ! The Porter deployment provider is removed; deployments now default to Kubernetes.
- ! Schema migrations run automatically on `helm upgrade` unless `<service>.database.migrations` is set to `false`; services that depend on new columns/tables (suspend/unsuspend, key deactivation, domain blocklist, Engine purge job) require the migrations to run.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 25, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue