THE CYBER TOOLCHAIN

NO. 133

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 7, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 142 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 33 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: Phoenix Security lets teams cap and review remediation load before it hits developers; Netcap can triage

vim ? help

offline with signed enrichment data; Prowler can enforce AWS policy from CloudTrail events instead of waiting for the next scan.

PROTECT

Throttle vulnerability work before it overwhelms engineering

Remediation Targets set scoped fix quotas per asset, and internal review can hold findings back before developers see them. Security teams can turn a raw backlog into planned remediation batches instead of dumping every finding into engineering queues at once.

Phoenix Security

PROTECT

Enforce cloud policy when AWS activity happens

Prowler added CloudTrail-driven real-time AWS detection and remediation with automated policy enforcement. That makes misconfiguration control less dependent on periodic scans: a risky cloud change can be caught and acted on as the event lands.

Prowler

IDENTIFY

Carry trusted enrichment into offline network triage

Netcap's signed community database system uses Ed25519 verification and rollback protection, with signed enrichment adding CISA KEV attribution and passive web-technology detection; `net triage` now runs offline. Analysts can preserve provenance for enrichment data and keep triage moving when the investigation environment is disconnected.

Netcap

PROTECT

Move code scanning decisions closer to the developer and the reviewer

Semgrep made Guardian Policies and local developer rules generally available, while Checkmarx exposed scanning, triage assistance, remediation assistance and exploitability ranking through Claude Marketplace. The useful shift is less handoff: policy, prioritization and fix guidance can appear where code is written or reviewed, not only after a central scan queue runs.

Semgrep · Checkmarx

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

xalgorix

Feedback ⓘ v4.6.182 NOTES ↗ RANK

Autonomous AI pentesting agents — real-time reconnaissance, vulnerability detection, and exploitation orchestration. Go + TypeScript.

> Xalgorix flipped several defaults toward bounded, deterministic scanning — bounded agent context, port-pinned scan boundaries with no subdomain enumeration, and epoch-based coverage convergence — while also fixing model-ID handling and base-URL fallback for gateway and multi-vendor LLM providers.

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

01 **Bounded working context by default** BREAKING 88

`XALGORIX_BOUNDED_CONTEXT` now defaults to `true`: full tool output is archived under the scan directory, old results are replaced in the agent conversation by compact retrieval stubs, and `read_tool_output` returns the original byte-identical content on demand. Tune retention with `XALGORIX_TOOL_ARCHIVE_ACTIVE_WINDOW` and `XALGORIX_TOOL_ARCHIVE_MIN_BYTES`, or disable with `XALGORIX_BOUNDED_CONTEXT=false`. Setups relying on the previous unbounded default, where full tool output was resent every iteration, will behave differently on upgrade.

Run a long autonomous scan without the agent context ballooning — old tool output is archived and swapped for stubs automatically under the new default.

```
$ XALGORIX_TOOL_ARCHIVE_ACTIVE_WINDOW=10
XALGORIX_TOOL_ARCHIVE_MIN_BYTES=4096 xalgorix --web
```

Opt out of bounded context for a short scan where you want all raw tool output kept live in the agent conversation at all times.

```
$ XALGORIX_BOUNDED_CONTEXT=false xalgorix --web
```

— Names all tuning env vars and gives opt-out and runnable examples v4.6.182

02 **Port-pinned scan boundaries** BREAKING 80

Scans are now limited to the port specified in the target URL — requests, port scans, and findings on any other port of that host are refused, so a scan of

`https://example.com:9000` stays on `:9000` only; service enumeration must name the pinned port, e.g. `-p 9000`. Scans that previously relied on automatic enumeration and coverage of additional ports on the same host will no longer produce those findings.

— Names flag and gives before/after but no runnable example v4.6.181

03 **Epoch-based coverage convergence** BREAKING 79

`XALGORIX_EPOCH_COVERAGE` now defaults to `true`, measuring coverage against a frozen surface snapshot that advances in bounded epochs so scans can close on large or still-growing attack surfaces without dropping any discovered endpoint. Set `XALGORIX_EPOCH_COVERAGE=false` to restore the previous behaviour.

Opt out of epoch-based coverage convergence when running a scan in an environment where the previous unbounded behaviour is required.

```
$ XALGORIX_EPOCH_COVERAGE=false xalgorix --web
```

— Explains mechanism and gives env-var opt-out command v4.6.181

04 **Model ID preservation for gateway LLM providers** IMPROVED 76

Gateway and inference-host providers named in `XALGORIX_LLM_PROVIDER` (OpenRouter, Groq, Together, vLLM, LM Studio, and others) now receive requests with the model ID unchanged — a slash-prefixed ID such as `google/gemini-3-pro` or `qwen/qwen3-32b` is no longer stripped before forwarding. Xalgorix also logs a one-time message when a provider setting overrides a model-ID prefix, making a stale `XALGORIX_LLM_PROVIDER` or `XALGORIX_API_BASE` setting easy to spot.

— Names config key and model IDs but no runnable command v4.6.180

05 **Catalog base URL fallback for multi-vendor providers** BREAKING 67

Built-in multi-vendor providers such as OpenRouter and Together now use their catalog base URL automatically when `XALGORIX_API_BASE` is unset, instead of falling back to a generic default host. Setups that relied on the previous generic fallback host will now send traffic to a different base URL.

— Names the config key and before/after but no example v4.6.180

THINNER COVERAGE BELOW

06 **No subdomain enumeration for single targets** BREAKING 52

Single-target scans no longer enumerate subdomains: subdomain discovery, takeover checks, and sibling subdomain content discovery are skipped, though a `www.` variant still counts as the target. Workflows that depended on automatic subdomain discovery and takeover checks in single-target mode will no longer trigger them.

— Clear before/after but no flag or command to act on `v4.6.181`

└──▷ BREAKING ON UPGRADE

! `XALGORIX_BOUNDED_CONTEXT` now defaults to `true`; any setup that depended on the previous default of `false` (unbounded context, full tool output resent every iteration) will behave differently on upgrade — set `XALGORIX_BOUNDED_CONTEXT=false` to restore the old behavior.

! `XALGORIX_EPOCH_COVERAGE` now defaults to `true`; any setup that relied on the previous off-by-default behaviour must set `XALGORIX_EPOCH_COVERAGE=false` explicitly to restore it.

! A port in the target URL now restricts the scan to that port only — scans that previously relied on automatic enumeration and coverage of additional ports on the same host will no longer produce those findings.

! Single-target scans of a registrable domain no longer enumerate or cover subdomains — workflows that depended on automatic subdomain discovery and takeover checks in single-target mode will no longer trigger them.

! Built-in multi-vendor providers (OpenRouter, Together, and others with no legacy default) now use their catalog base URL when `XALGORIX_API_BASE` is unset — setups that relied on the previous generic fallback host will now send traffic to a different base URL.

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

Datadog pathrunner

Feedback ⓘ v1.1.0 NOTES ↗ RANK

> pathrunner v1.1.0 adds new AWS identity loading options and a guided discover wizard for module configuration.

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

01 **Identity loading from env vars and IMDS** NEW 80

The `identity add` command gains `--from-env-vars` and `--from-imds` flags, letting pathrunner load AWS identity from environment variables (e.g. inside a CI job or assumed-role shell) or from EC2/Lambda instance metadata, authenticating automatically as the execution role when run on an AWS-hosted instance.

Pick up credentials already present in the environment (e.g. inside a CI job or an assumed-role shell) without specifying a profile.

```
$ pathrunner> identity add --from-env-vars
```

Authenticate as the EC2 or Lambda execution role automatically when running pathrunner on an AWS-hosted instance.

```
$ pathrunner> identity add --from-imds
```

– Names exact flags and command with runnable examples v1.1.0

THINNER COVERAGE BELOW

02 **Tiered discover wizard in REPL** NEW 35

pathrunner's REPL now includes a tiered discover wizard that auto-discovers required module options and prompts interactively for optional ones.

– Describes behaviour but no command or UI path given v1.1.0

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

taranis-ai Feedback ⓘ 1.5.3 NOTES ↗ RANK

Taranis AI is an advanced Open-Source Intelligence (OSINT) tool, leveraging Artificial Intelligence to revolutionize information gathering and situational analysis.

> Taranis AI's 1.5.3 release adds live SSE-based updates to the Scheduler and My Tasks views and surfaces report types in the Assess add-to-report dialog.

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

01 Live SSE updates for Scheduler and My Tasks IMPROVED

50

The Scheduler and My Tasks views now update live over SSE, so task and queue changes appear without a page reload, while existing filters and pagination state are preserved.

– Explains mechanism and behaviour but no config or flag 1.5.3

02 Report type shown in Assess add-to-report dialog IMPROVED

35

The Assess add-to-report dialog now shows the report type when adding stories to an existing report.

– Clear UI change but minimal detail on scope 1.5.3

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

Hak5 Shark Jack Display Feedback ⓘ

sold out · Nov 6, 2026 STORE ↗

Jack into a network and instantly run advanced recon, exfiltration, attack and automation payloads.

Incognito Black, Pro Bundle (sjd-incognito-pro) sold out at `shop.hak5.org` on 6 Nov 2026.

WAS THIS USEFUL? 👍 👎

◆ Code Security

Semgrep

Feedback ⓘ updated · Nov 7, 2026 NOTES ↗ RANK

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

> Semgrep's latest release brings Guardian Policies and local developer rules to general availability, adds experimental Objective-C scanning, and expands CSV exports and API fields for triage workflows, alongside parser improvements for C#, Terraform, and Lua.

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

01 Guardian Policies and local developer rules GA NEW

95

Semgrep Guardian Policies and local developer rules are now generally available for all organizations. Developers can store local rules in `~/.semgrep/guardian-rules/` or the root of a repository, with Guardian running up to 1,000 organization rules and 50 local developer rules per scan, and per-rule control available in the Claude Code remote plugin and the Codex plugin.

Drop custom Guardian rules into the per-user directory so they run locally in your AI coding agent without affecting CI policy.

```
$ cp my-org-rules.yaml ~/.semgrep/guardian-rules/
```

— Names exact paths, numeric limits, and plugin integrations snapshot-20261107

02 Language parsing improvements for C#, Terraform, and Lua

IMPROVED

65

Semgrep Code now parses C# Unicode escape sequences and Unicode identifier characters, Terraform provider functions with a namespace (e.g. provider::fn(...)), conditional and loop template directives, and Lua 5.4 variable attributes plus Lua 5.5 leading attributes, global declarations, and named variable-argument parameters.

— Many named syntax features but nothing for a reader to run snapshot-20261107

03 Findings CSV export gains triage and ticket columns IMPROVED

60

Findings CSV exports now include Last triaged by, Ticket key, and Ticket URL columns; the ticket columns appear only when a Jira integration is active.

— Named columns and condition, but simple addition snapshot-20261107

04 Managed Scans source restrictions for Free/Teams plans

60

BREAKING

Free and Teams plan organizations can no longer select GitHub Enterprise, GitLab Self-Managed, or Bitbucket Data Center when setting up Semgrep Managed Scans.

— Names affected plans and source types but no workaround snapshot-20261107

THINNER COVERAGE BELOW

05 Advisories-with-findings filter NEW

55

A new Advisories with findings filter on the Advisories page shows only advisories affecting your projects, respecting project-level role-based access control and excluding archived projects.

— Clear UI location but no deeper mechanism detail snapshot-20261107

06 Experimental Objective-C scanning NEW

55

Semgrep adds experimental support for Objective-C, parsing `.m` files and running search-mode and taint-mode rules with metavariables and ellipsis operators.

— Describes scope but marked experimental with no usage example snapshot-20261107

07 Self-service Semgrep Secrets purchasing NEW

50

Self-service organizations can now buy Semgrep Secrets standalone or bundled, and manage or cancel subscriptions via the Stripe customer portal by clicking Edit plan.

— Names UI action but little mechanism beyond that snapshot-20261107

08 Secrets bulk triage now requires triage permission BREAKING

50

Users with the Read-only role and API tokens with the Read-only scope can no longer bulk ignore or reopen Semgrep Secrets findings; bulk triage now requires triage permission.

— Clear before/after but no migration steps given snapshot-20261107

09 issueUrl field in issues API IMPROVED

35

The issues API now returns an `issueUrl` field containing a link to each issue in AppSec Platform.

— Names the field but no endpoint path or usage shown snapshot-20261107

└──▷ BREAKING ON UPGRADE

! Users with the Read-only role and API tokens with the Read-only scope can no longer bulk ignore or reopen Semgrep Secrets findings; bulk triage now requires triage permission.

! Free and Teams plan organizations can no longer select GitHub Enterprise, GitLab Self-Managed, or Bitbucket Data Center when setting up Semgrep Managed Scans.

WAS THIS USEFUL? 👍 👎

psalm-plugin-laravel

Feedback ⓘ v4.19.0 NOTES ↗ RANK

Security-focused Laravel static analyser

> psalm-plugin-laravel v4.19.0 turns on four previously-opt-in runtime checks by default and adds type-aware support for Livewire component views, Eloquent factory relation helpers, and the Laravel bootstrap entrypoint.

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

01 **Four runtime checks enabled by default** BREAKING 85

`MissingView`, `UnconfiguredFilesystemDisk`, `SerializedQueuedModel`, and `UnknownModelAttribute` now run by default at `errorLevel` 1–4 (reported as `info` above that level). Projects that still have `<findMissingViews value="false" />` in `psalm.xml` (added automatically by `psalm-laravel init` v3.11.0+) must delete that line to actually get the newly default `MissingView` check.

Remove the now-redundant opt-out so the newly default `MissingView` check is active in your project.

```
$ # In psalm.xml, delete this line:
# <findMissingViews value="false" />
# Then run:
./vendor/bin/psalm-laravel analyze
```

— Names all four checks, error levels, and exact config line to remove v4.19.0

02 **Typed Eloquent factory relation helpers** NEW 65

psalm-plugin-laravel now types has{Relation}() and for{Relation}() magic methods on Eloquent factories, eliminating false `UndefinedMagicMethod` errors previously raised on calls like User::factory()->hasLinks(2)->create().

Catch a misspelled Eloquent factory relation helper (e.g. `haslinks`) that previously raised `UndefinedMagicMethod` — it now type-checks cleanly.

```php
$user = User::factory()->hasLinks(2)->create();
// $user is now typed as User — no UndefinedMagicMethod error
```

— Names the magic methods and shows the fixed error with example code v4.19.0

THINNER COVERAGE BELOW

03 **Typed Livewire component views** NEW 40

psalm-plugin-laravel now types Livewire component views from their class, resolving `$this` and public properties inside the view templates.

— Explains behaviour but gives no command or example v4.19.0

04 **Typed bootstrap entrypoint instance** NEW 40

psalm-plugin-laravel now types `require 'bootstrap/app.php'` in `artisan` and `public/index.php` as the booted `Application` instance.

— Names the files and type but no runnable example v4.19.0

└──▷ **BREAKING ON UPGRADE**

! Projects generated with `psalm-laravel init` v3.11.0+ have `<findMissingViews value="false" />` in `psalm.xml`; without removing it, the newly default `MissingView` check remains suppressed.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Copilot code review: API management and review overview updates (2026-11-07)

BLOG ↗ RANK

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

> GitHub shipped general availability of @GitHub in Slack and Teams with Canvas-based review, new REST/GraphQL APIs for managing and attributing Copilot code review across enterprise, organization, and repository levels, broader usage tracking in the Copilot impact dashboard, and a wave of agent improvements across VS Code, Copilot CLI, and JetBrains Air.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 **Code review REST API enhancements** NEW 90

The REST API now exposes explicit attribution fields identifying Copilot's reviews and comments — including the review effort used and the skills, MCP servers, and custom instructions reported as sources for a comment — and adds code review settings endpoints for enterprise, organization, and repository levels: enterprise administrators read and update code review enablement across organizations, organization administrators manage default review effort and runner configuration, and repository administrators manage review effort, custom instructions, and MCP tool settings, with repository settings endpoints supporting GitHub Apps and fine-grained personal access tokens via `Administration: read` and `Administration: write` permissions.

— Names attribution fields but not the exact endpoint path Names exact permission scopes and admin levels, though not literal endpoint paths

Copilot code review: API management and…

02 **Chat and VS Code agent usage tracking in impact dashboard**

NEW 85

The Copilot impact dashboard's Agent adoption chart now shows Copilot Chat and the VS Code Agents window, bringing the total tracked experiences to nine. Enterprise and organization 28-day reports expose this via new `chat` and `vscode_agent` entries in `copilot_feature_engagement.totals_by_feature`.

— Names exact API fields and chart scope Copilot impact dashboard adds chat and …

03 **GitHub Copilot in Slack and Teams reaches general availability**

NEW 85

@GitHub is now generally available in Slack and Teams, letting users investigate bugs, create issues, implement changes, and open pull requests directly from chat conversations. In Teams, Canvas lets you review implementation plans, code changes, and previews alongside team discussions without leaving the app. @GitHub automatically identifies the repository from your message and launches a coding session without manual repository selection, and issues and pull requests it creates include 'Created from a Slack/Teams conversation with @username' in their description.

▶ Watch video

Start a coding session from Slack or Teams to investigate a bug and open a pull request without switching tools.

📍 Steps: 1. In your Slack or Teams channel, type @GitHub followed by your request, e.g. 'investigate the login timeout bug in this repo and open a PR with a fix'. 2. @GitHub identifies the relevant repository from your message and launches a coding session automatically. 3. Review the implementation plan, code changes, and preview in Canvas (Teams) or the conversation thread (Slack). 4. Confirm to create the pull request — it will include 'Created from a Slack/Teams conversation with @username' in its description.

— Includes a full runnable workflow example and named attribution text

@GitHub in Slack and Teams is generally…

---

04 **/settings dashboard in GitHub Copilot CLI** NEW 75

Manage all CLI plugins, MCP servers, and settings from one place using the `/settings` dashboard in GitHub Copilot CLI.

Open all Copilot CLI plugin, MCP server, and general settings from a single searchable dashboard.

```
$ /settings
```

— Exact runnable CLI command provided GitHub Copilot weekly releases — Novemb…

---

05 **Cron-scheduled local automations in VS Code** NEW 70

Run local automations on a schedule by adding cron expressions and custom agents from your workspace in VS Code.

Schedule a nightly security scan or build task without leaving VS Code by wiring a cron expression to a custom agent.

📍 In VS Code, open Local Automations, add a new automation, set the schedule to a cron expression (e.g. '0 2 * * *'), and select a custom agent from your

workspace.

— Includes a runnable cron example and workspace context

GitHub Copilot weekly releases — Novemb…

06 **Attach external files to chat without workspace copy** NEW

70

Add files to chat without moving them into your workspace using `Add Context > Files` in VS Code; cloud sessions support images only.

Attach a mockup screenshot to a Copilot chat in VS Code without copying it into the project directory.

📍 In the VS Code chat input, click Add Context › Files and select the screenshot or mockup file from anywhere on disk.

— Names exact menu path and a scope limitation GitHub Copilot weekly releases — Novemb…

07 **Review overview layout redesign** IMPROVED

65

Review overview now shares one Files not reviewed section for skipped and content-excluded files with filenames and linked exclusion reasons, omitting empty sections and showing an overflow count for long lists. Effort used and surveys moved into a compact footer so findings stay together at the top, and findings unrelated to recent changes collapse then hide on later reviews of the same pull request.

— Describes concrete UI behaviour change but no direct link or command

Copilot code review: API management and…

08 **Structured feedback reasons for review overview** NEW

65

After giving a review overview a thumbs-down, use Add a reason to pick the main problem — such as a disagreeable approval assessment, a verbose comment, or incorrect or nitpicky findings — to send structured feedback to improve Copilot's findings.

— Names the exact UI action and feedback categories

Copilot code review: API management and…

09 **Bypass list support for automatic review rule** IMPROVED

60

The Automatically request Copilot code review repository rule now supports the bypass list, so exempt actors will not receive an automatic review.

— Names the exact repository rule and its new behaviour

Copilot code review: API management and…

10 **GitHub Copilot support in JetBrains Air** NEW

60

GitHub Copilot now works in JetBrains Air, letting you track multiple agent sessions and run tasks locally or in the cloud through Air Teams with an existing Copilot subscription.

— Names the IDE and Air Teams but no setup steps GitHub Copilot weekly releases — Novemb…

THINNER COVERAGE BELOW

11 **Remote previews and interactive canvases in VS Code Agents window**

NEW

55

Adds remote browser previews, remote app previews, and interactive canvases in the VS Code Agents window, enabled via the remote browser support setting.

.

— Names the setting but gives no exact steps to enable it

GitHub Copilot weekly releases — Novemb…

12 **Agent Merge setup from pull request** NEW

55

Set up Agent Merge directly from a pull request so Copilot merges it automatically when required checks and rules pass.

— Describes trigger conditions but no exact UI path

GitHub Copilot weekly releases — Novemb…

13 **Assisted approval personal instructions for tool requests**

NEW

50

Give Copilot standing guidance on which tool requests to approve without prompting, via opt-in personal instructions for assisted approval that apply to local sessions.

— Explains mechanism generally but no exact config surface named

GitHub Copilot weekly releases — Novemb…

14 **Context limit compaction and usage review** IMPROVED

45

Keep working when a conversation hits its context limit by compacting the conversation or reviewing context usage directly from the error.

— Describes behaviour but no exact command or UI path

GitHub Copilot weekly releases — Novemb…

15 **Start Copilot app session by dropping files** NEW 45

Start a Copilot app session by dropping files onto New in the sidebar.

— Exact UI interaction named but minimal context GitHub Copilot weekly releases — Novemb…

16 **'Passive users' cohort renamed to 'Chat-first users'** IMPROVED 35

The Passive users adoption cohort in the Copilot impact dashboard is now named Chat-first users to better describe users who are active in Copilot but do not qualify for an adoption phase.

— Simple rename with no further mechanism Copilot impact dashboard adds chat and …

17 **AI credit usage shown in review agent sessions** NEW 30

Copilot code review agent sessions now show the AI credits consumed in a review.

.

— Bare statement with no further mechanism or numbers

Copilot code review: API management and…

└──▷ ALSO FROM THESE RELEASES

▶ Watch video

▶ Watch video

WAS THIS USEFUL? 👍 👎

Checkmarx

Feedback ⓘ

Checkmarx Joins Claude Marketplace to Advance AI-Powered Application Security (2026-11-0

BLOG ↗ RANK

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

> Checkmarx brought its Fusion SAST, Triage Assist, and Remediation Assist products to the Claude Marketplace, letting Anthropic customers in the U.S. pair Checkmarx scanning and exploitability ranking with Claude AI.

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

01 **Fusion SAST available in Claude Marketplace** NEW 45

Checkmarx Fusion SAST now pairs Checkmarx scanning with Claude AI to surface real vulnerabilities in shipped code, available to Anthropic customers in the U.S. through the Claude Marketplace.

— States mechanism but no concrete config or command to act on

Checkmarx Joins Claude Marketplace to A…

02 **Triage Assist exploitability ranking** NEW 40

Triage Assist ranks findings by attacker exploitability so security teams review the most critical issues first, now reachable via the Claude Marketplace integration.

— Explains ranking behaviour but lacks specifics like scoring method

Checkmarx Joins Claude Marketplace to A…

03 **Remediation Assist added to Claude Marketplace** NEW 15

Checkmarx Remediation Assist was added alongside Fusion SAST and Triage Assist to the Claude Marketplace for Anthropic customers in the U.S., though no functional detail beyond its availability was given.

— Only named, no description of what it does Checkmarx Joins Claude Marketplace to A…

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Package Firewall updates (2026-11-06) BLOG ↗

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

Veracode Package Firewall is now available, automatically blocking untrusted packages across artifact repositories and package managers to enforce dependency compliance.

NEW DETECTION Proactive Maven package monitoring for malware analysis

NEW SUPPORTED PLATFORM Package Firewall available in the European region

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

plumber

Feedback ⓘ v0.7.23 NOTES ↗ RANK

Plumber detects CI/CD security issues in your GitHub workflows and gives you a score

> Plumber's platform banner now surfaces when snapshot lanes carry over from a previous pass.

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

01 **Carry-over snapshot lanes in platform banner** IMPROVED

35

The platform banner now names the snapshot lanes carried from an earlier pass, making carry-over state visible at a glance.

– Describes UI change but no mechanism or exact location v0.7.23

WAS THIS USEFUL? 👍 👎

Sigstore cosign

Feedback ⓘ CVE-2022-23649 CVE ↗

Code signing and transparency for containers and binaries

`CVE-2022-23649`: Cosign can be manipulated to falsely claim an entry exists in the Rekor transparency log (affects `< 1.5.1`)

WAS THIS USEFUL? 👍 👎

◆ Application Security Posture

Phoenix Security

Feedback ⓘ

Remediation Targets – Configuration & Usage Guide (2026-11-06) BLOG ↗ RANK

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

> Phoenix Security introduces Remediation Targets, letting security teams set scoped per-asset fix quotas for findings and hold back findings pending internal review before they reach development teams.

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

01 **Custom status workflow to hold back findings** NEW 79

Teams can create a custom status via Settings > Organization > Configuration > Custom Status Configuration > Add Status, apply it to findings in Vulnerabilities > Findings, then add an Add Exclusions (BUT NOT) rule on the Objective filtered by Finding Status to keep those findings from counting toward a team's remediation target until review is complete. Held-back findings can be located again via the Custom Statuses filter in the Findings Filters panel, or shown/hidden using the Show Excluded options toggle.

Hold back a cluster of findings pending security review so they never surface to the development team until cleared.

📍 1. Go to Settings › Organization › Configuration, expand Custom Status Configuration, select Add Status, enter Security review, pick a colour, and select Confirm. 2. In Vulnerabilities › Findings, select the findings to hold back, choose Update Statuses, and apply Security review. 3. Open the Objective, select Add Exclusions (BUT NOT), set category to Finding Status, choose Is any of, and select Security review. 4. Select Save target, then Re-run selection on the Remediation Targets page.

Review all findings currently held under a custom status to decide when they are ready to be released to the team.

📍 1. Go to Vulnerabilities › Findings and select Filters. 2. Under Custom Statuses, choose Security review and select Apply Filters. 3. When guidance is ready, remove the status from the findings (or remove the exclusion from the Objective) and select Re-run selection.

– Exact menu paths and filter names given for the whole workflow.

Remediation Targets – Configuration & U…

02 **Scoped remediation Objectives with per-asset quotas** NEW

77

Phoenix Security adds Remediation Targets, where an Objective can be scoped to the organisation, or to a specific application, environment, component or service using the Scopes card, with the nearest assignment winning. Objectives match findings by category (e.g. Risk Level set to High/Critical) and set a goal such as 'Allow up to 5 open matching findings per asset', run in `Objective-based` mode via `Re-run selection` under Automation > Remediation Targets.

Create a Remediation Target that asks each team to fix all Critical and High findings beyond a per-asset allowance of 5.

> 📍 1. Go to Automation › Remediation Targets and select `New Objective`. 2. Name it `Critical and High`. In the MATCHING row set category to Risk Level and select `High` and `Critical`. 3. Set the goal to 'Allow up to 5 open matching findings per asset'. 4. Select `Save target`, open the Scopes card and assign it to the Organisation. 5. Set the mode to `Objective-based`, then select `Re-run selection` and wait for `Last run` to show Succeeded.

— Names exact UI path, goal syntax and scoping mechanism.

Remediation Targets – Configuration & U…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Mobile App Security

MobSF

Feedback ⓘ v4.6.4 NOTES ↗ RANK

Automated mobile application pen-testing, malware analysis and security assessment framework for Android, iOS and Windows apps, with static and dynamic analysis.

> Across v4.6.2 and v4.6.4, MobSF added direct `.dex` static analysis and a built-in `mitmproxy` `-backed` live Web Traffic viewer with fuzzer replay, shipped physical-device Android dynamic analysis with a Frida Gadget APK builder, and rounded out the release with dark mode, Flutter/split-APK static analysis improvements, a reworked PDF report pipeline, a revised security score formula and an expanded permissions knowledge base.

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

01 **Android dynamic analysis on physical devices with Frida Gadget**

NEW 85

Adds Android dynamic analysis support for physical devices, covering both rooted and stock (non-rooted) devices with hardened defaults for physical-device deployments. Includes a Frida Gadget APK builder to instrument apps on stock devices (including split APK sets), a gallery of installed apps on the connected device with on-demand scanning, and a live step-by-step progress display while the dynamic analysis setup runs.

— Names concrete tooling (Frida Gadget) but no exact command shown. v4.6.2

02 **Static analysis improvements for Flutter and split APKs**

IMPROVED 75

Static analysis adds Dart string extraction from Flutter `libapp.so`, a detection check that flags unauthenticated `LocalServerSocket` listeners, and analysis of native libraries inside split APKs. False-positive noise is reduced by skipping `JetBrains` and Flutter embedding classes and by ignoring Flutter plugin constants and tightening binary secret detection.

— Names specific checks and files but no CLI/API example. v4.6.2

03 **Built-in mitmproxy Web Traffic viewer with fuzzer replay** NEW

70

Dynamic analysis now includes a built-in authenticated live Web Traffic viewer backed by `mitmproxy`, with a `Send to Fuzzer` action that replays captured requests through an upstream proxy.

— Mechanism named but no UI path or command shown. v4.6.4

04 **PDF report generation rebuilt on xhtml2pdf** IMPROVED

65

PDF report generation now uses the pure-Python `xhtml2pdf` library with a `ResourceAccessPolicy` for confined asset resolution, replacing the prior `wkhtmltopdf` and `pdfkit` dependencies.

— Internal dependency swap, no user-facing action required. v4.6.4

THINNER COVERAGE BELOW

05 **Direct static analysis of standalone .dex files** NEW

50

MobSF can analyze standalone Android `.dex` files directly with static analysis, without needing to wrap them in an APK first.

— Clear capability named but no workflow example shown. v4.6.4

06 **Expanded Android permissions knowledge base** IMPROVED

40

Adds 26 new Android permissions to the Android permissions knowledge base used during static analysis.

— Just a count added, no detail on which permissions or effect. v4.6.4

07 **Dark mode across admin shell and traffic viewer** NEW

30

MobSF adds a dark mode toggle to the AdminLTE admin shell, and the new live Web Traffic viewer also ships with dark mode support.

— UI toggle described without navigation path. v4.6.4 v4.6.2

08 **Baseline-anchored security score formula** IMPROVED 30

MobSF's AppSec security score calculation now uses a baseline-anchored monotonic formula.

— No explanation of how the new formula computes scores. v4.6.4

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

CHIPSEC

Feedback ⓘ 2.0.10 NOTES ↗ CODE ↗ RANK

Platform Security Assessment Framework

> CHIPSEC 2.0.10 adds an offline Secure Boot DBX revocation check and optional image-hash normalization for UEFI binary scanning.

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

01 **Rebase-0 SHA-256 normalization in scan_image** IMPROVED

50

CHIPSEC adds optional rebase-0 SHA-256 normalization for PE32 and TE images in `tools.uefi.scan_image`.

— Names module path but not the flag to enable it 2.0.10

02 **Offline Secure Boot DBX check utility** NEW

37

CHIPSEC adds an offline Secure Boot DBX check utility to verify revoked signatures without requiring a live system.

— Names capability but no command or flag given 2.0.10

WAS THIS USEFUL? 👍 👎

ntoseye

Feedback ⓘ v0.53.0 NOTES ↗ RANK

WinDbg-like kernel debugger for Windows, from Linux and macOS

> ntoseye v0.53.0 adds minidump handle inspection, typed globals from private PDBs, per-processor selection on the memory backend, and process-scoped breakpoints, alongside five breaking changes that consolidate trace environment variables, restructure the SDK's PteWalk result, tighten dt/dl limits, realign sx event aliases, and change two environment-variable boolean semantics.

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

01 Process-scoped breakpoints and watchpoints NEW

90

Scope breakpoints to a process with `/p <pid>` on `bp`, `ba` and `bm` — for example `ba e1 /p <pid> kernelbase!GetFileAttributesExW+0x60` — without changing the current process context; the SDK's breakpoints.add(hardware=True, process=...) and breakpoints.watch(..., process=...) provide the same capability.

Set a hardware watchpoint scoped to a specific process without switching process context — handy when you are stopped in the System process or the hypervisor and want to catch an access in a user process.

```
$ ba e1 /p 1234 kernelbase!GetFileAttributesExW+0x60
```

— Runnable example with exact SDK calls v0.53.0

02 Unified NTOSEYE_TRACE diagnostic variable BREAKING

90

A single `NTOSEYE_TRACE` environment variable, set to a comma-separated list of `kd`, `kd-bytes`, `gdb`, `step` and `unwind`, now controls all diagnostic traces, replacing the five separate `NTOSEYE_KD_TRACE`, `NTOSEYE_KD_TRACE_BYTES`, `NTOSEYE_GDB_TRACE`, `NTOSEYE_STEP_TRACE` and `NTOSEYE_UNWIND_TRACE` variables, which are removed.

Enable only KD-protocol and unwind tracing for a session to diagnose stop-time problems without the noise of GDB or step traces.

```
$ NTOSEYE_TRACE=kd,unwind ntoseye <target>
```

— Runnable example plus full before/after variable list v0.53.0

03 **Handle object inspection in minidump analysis** NEW

78

ntoseye can inspect handle object information from a user-mode minidump with `!handle <handle> f`, surfacing thread/process IDs, TEB/PEB, priority, exit status, event/mutant/semaphore state, and section size via `UserHandle.object_info` in the SDK.

– Command named but no worked example v0.53.0

04 **sx event alias additions and realignment** BREAKING

75

ntoseye adds `bpec`, `ssec`, `cce`, `ch`, `hc`, `wkd`, `aph`, `dbce`, `vcpp`, `clr` and `clrn` as `sx` event aliases, and realigns existing aliases to WinDbg's event table: `cc` now maps to Ctrl+C (0x40010005) instead of the C++ exception (previously `eh`), `asrt` is now an assertion failure (0xC0000420), and `bp` is now the breakpoint exception alias `bpe`.

– Full alias table but no worked example v0.53.0

05 **Processor selection on memory backend** NEW

73

Select processor N on the `memory` backend with `~N` and `~Ns`, so `!idt`, `!gdt`, `!pcr` and PRCB-reading commands inspect that processor as on a halted target.

– Exact command syntax but no worked example v0.53.0

06 **PteWalk result restructured into levels array** BREAKING

73

The SDK's `PteWalk` result now lists page-table entries in a `levels` array (root first), replacing the individual `pxe`, `ppe`, `pde` and `pte` fields.

– Exact field rename given, no migration example v0.53.0

07 **Typed globals from private PDBs via dx/dt** IMPROVED

70

`dx` and `dt` now show typed globals and static variables from private PDBs — for example `dx balloon!GUID_DEVINTERFACE_BALLOON` reads the GUID's fields without requiring a manual cast.

Read typed fields of a private-PDB global directly — useful when debugging your own driver without having to cast the symbol manually.

```
$ dx balloon!GUID_DEVINTERFACE_BALLOON
```

– Runnable example but limited mechanism detail v0.53.0

08 Environment-variable boolean semantics standardized

BREAKING

68

`NTOSEYE_NO_PDB_FROM_MEMORY` and `NTOSEYE_GDB_ON_HVF` are now off when empty or `0` and on for any other value; previously `NTOSEYE_NO_PDB_FROM_MEMORY=0` disabled the feature and `NTOSEYE_GDB_ON_HVF` was on only at `1`.

— Exact before/after values but no example `v0.53.0`

---

09 dt -a and dl size limits now enforced as errors

BREAKING

60

`dt -a` counts above 16 and `dl` sizes above 64 are now errors instead of being silently truncated.

— Clear before/after but no example `v0.53.0`

└──▷ BREAKING ON UPGRADE

! The SDK's `PteWalk` result now lists page-table entries in a `levels` array (root first), replacing the individual `pxe`, `ppe`, `pde` and `pte` fields.

! The `sx` alias `cc` now maps to Ctrl+C (0x40010005) instead of the C++ exception (previously `eh`); `asrt` is now an assertion failure (0xC0000420); `bp` is now the breakpoint exception alias `bpe`.

! `dt -a` counts above 16 and `dl` sizes above 64 are now errors instead of being silently truncated.

! `NTOSEYE_TRACE` replaces `NTOSEYE_KD_TRACE`, `NTOSEYE_KD_TRACE_BYTES`, `NTOSEYE_GDB_TRACE`, `NTOSEYE_STEP_TRACE` and `NTOSEYE_UNWIND_TRACE`; the old variables are removed.

! `NTOSEYE_NO_PDB_FROM_MEMORY` and `NTOSEYE_GDB_ON_HVF` are now off when empty or `0` and on for any other value; previously `NTOSEYE_NO_PDB_FROM_MEMORY=0` disabled the feature and `NTOSEYE_GDB_ON_HVF` was on only at `1`.

WAS THIS USEFUL? 👍 👎

AsmResolver

Feedback ⓘ v6.1.0 NOTES ↗ RANK

A library for creating, reading and editing PE files and .NET modules.

> AsmResolver v6.1.0 adds PE load configuration data directory support with safer relocation handling, improves .NET Core runtime path resolution, and speeds up type lookups in larger .NET binaries.

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

01 **DOTNET_ROOT preference in DotNetCorePathProvider** IMPROVED

60

`DotNetCorePathProvider` now prefers `DOTNET_ROOT` installations of the runtime/SDK to better reflect loader behavior.

– Names exact class and environment variable involved v6.1.0

02 **Safe binary reading with TryReadX methods** NEW

60

Added `TryReadX` methods to `BinaryStreamReader` for safer binary parsing.

– Names exact method pattern and class but no usage example v6.1.0

THINNER COVERAGE BELOW

03 **PE load configuration data directory support** NEW

45

AsmResolver can now read and write PE load configuration data directories. Base relocations are deduplicated by default when reassembling PEs.

– Names the feature and a behaviour change but no API names v6.1.0

04 **Faster type lookup in large .NET binaries** IMPROVED

30

Improved type lookup performance by namespace and name for larger .NET binaries.

– States the improvement but no mechanism or benchmark given v6.1.0

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

mac_apt

Feedback ⓘ v.1.39.0 NOTES ↗

macOS (& ios) Artifact Parsing Tool

mac_apt adds a SIRI plugin for Siri history, a LESSHST plugin for less pager history, and extends TERMSESSIONS to read .python_history.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v.1.39.0 https://github.com/ydkhatri/mac_apt.git
# already have the repo? check out this version:
$ git checkout v.1.39.0
```

New `SIRI` plugin reads Siri history from macOS artifacts.

New `LESSHST` plugin reads less pager history (search patterns, shell commands, and marks) from `.lesshst` for every user.

The `TERMSESSIONS` plugin now also reads `.python_history` alongside bash and zsh history.

NEW INTEGRATION Island browser parsing in CHROMIUM plugin

NEW SUPPORTED PLATFORM aarch64 syscall support

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.15.3 NOTES ↗ CODE ↗ RANK

A framework for secure and scalable network traffic analysis - https://netcap.io

> Netcap shipped a signed community database system with Ed25519 verification and rollback protection, a new signed enrichment channel adding CISA KEV attribution and passive web-technology detection, expanded evidence-recovery protocol reconstruction, and an offline `net triage` command — alongside breaking changes retiring legacy database endpoints and raising the macOS requirement for DPI.

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

01 **Signed community database system** NEW 95

Install the community database bundle with `net util -download-dbs` or refresh it with `net util -update-dbs`; the layout-3 database (NVD 2002–2026 schema 2, DB-IP City/ASN, IANA services, UA-parser, and HASSH from a pinned BSD-licensed source) is cryptographically signed with Ed25519 and protected against rollback, rejecting expiring updates, future timestamps, lower key epochs or sequences, and same-sequence equivocation while staying usable offline. The database server supports `keys`, `approve`, `sign`, `serve`, `candidate`, and `pack` modes via `-mode`, with flags `-policy`, `-policy-sha256`, `-previous-count`, `-key`, `-certificate`, `-sequence`, `-epoch`, and `-signed-root` (also readable from `NC_DBS_SIGNED_ROOT`); install-time trust can be overridden for private publishers via `NC_DBS_APPROVALS` and `NC_DBS_APPROVALS_SHA256`. Windows database activation now works without symlink privileges.

Download the signed community databases for the first time on a fresh Netcap installation.

```
$ net util -download-dbs
```

Refresh an existing database installation to pull the latest reviewed and signed snapshot from the community publisher.

```
$ net util -update-dbs
```

Point Netcap at a private approved-source registry and independently pin its SHA-256 for air-gapped or enterprise deployments.

```
$ export NC_DBS_APPROVALS=/path/to/private-approvals.json
export NC_DBS_APPROVALS_SHA256=<sha256-of-approvals-file>
net util -update-dbs
```

First-time install of the community database bundle, including NVD feeds and GeoIP data, without any extra configuration.

```
$ net util -download-dbs
```

Upgrade an existing database installation to the latest reviewed snapshot.

```
$ net util -update-dbs
```

— Exact commands, env vars and verification mechanism all given. v0.15.2 v0.14.0

02

Signed enrichment channel: CISA KEV, web-tech detection, device catalogues

NEW

90

Download the signed optional enrichment channel with `net util --download-enrichment`; it installs under `NC_ENRICHMENT_ROOT` (default `$NC_CONFIG_ROOT/enrichment`) using bounded member verification and monotonic sequence checks, and the approval chain can be overridden via `NC_DBS_APPROVALS` and `NC_DBS_APPROVALS_SHA256` for private publishers. It adds CISA KEV attributed CVE enrichment with source links (federal deadlines labeled as such, Unknown ransomware labels not treated as negative assertions), passive web technology detection with bounded origin-transaction advertisements covering WordPress/Drupal built-ins and reviewed optional Hugo/Joomla descriptors (no active probes or inferred exact versions from major-only headers), optional historical FingerBank option-55 device candidates under ODbL/DbCL and IEEE vendor prefixes (falling back to the bounded MAC fallback when absent), and popularity filtering controlled by an explicit external CSV only, off by default, with PSL-aware filtered Maltego transforms and no bundled popularity dataset. Attribution metadata can be retrieved via `curl http://127.0.0.1:8080/api/enrichment/attribution`.

Bootstrap the signed optional enrichment channel on a new host to gain CISA KEV, device catalogue, and passive web-tech detection data.

```
$ net util --download-enrichment
```

Store enrichment data in a custom path instead of the default location under `$NC_CONFIG_ROOT/enrichment`.

```
$ NC_ENRICHMENT_ROOT=/data/netcap/enrichment net util --download-enrichment
```

Retrieve enrichment attribution metadata and licence notices programmatically to verify provenance before ingesting enriched records.

```
$ curl http://127.0.0.1:8080/api/enrichment/attribution
```

— Exact command, env vars and an API endpoint given. v0.15.3

---

03 **Legacy database endpoints and schema retired** BREAKING

70

Legacy v1/v2 database distribution endpoints are retired and now return HTTP 410, so installations using those endpoints must switch to `net util -update-dbs`. Schema-1 database readers must upgrade before consuming the new layout-3 database, and legacy CMS and DHCP OS catalogues are disabled, with HASSH catalogue enrichment unavailable in this database snapshot.

— Exact retired endpoints and a migration path are named. v0.14.0

---

04 **File capture no longer auto-starts Web UI** BREAKING

70

File capture now exits without starting the Web UI by default when `--read` is set; pass `--http <addr>` explicitly to opt back into the old behavior.

— Exact flag to restore prior behavior is given. v0.14.0

---

05 **`net triage` command for offline audit summaries** NEW

60

A new `net triage` command summarizes existing audit output from offline captures without re-running capture, usable as `net capture -read traffic.pcap -out nc && net triage --read nc`.

Analyze a PCAP offline and get a host/account summary without launching the Web UI, then optionally view full JSON output with provenance.

```
$ net capture -read traffic.pcap -out nc && net triage --read nc
```

— Runnable command shown directly in the example. v0.14.0

06 **Expanded evidence recovery and protocol reconstruction** NEW

60

Evidence recovery now reconstructs SMTP, IMAP, POP3, LDAP, and NTLM authentication conversations, handles request-only HTTP and SMB captures and post-gap sessions with request/response attribution and SMB file reassembly, reassembles fragmented TLS and QUIC hellos and QUIC Initials while recognizing mail protocol TLS upgrades, and correlates host identity across DHCP, NetBIOS, LDAP, Kerberos, and DCE/RPC observations without assigning identities to unrelated hosts.

— Many protocols named but no invocation surface given. v0.14.0

THINNER COVERAGE BELOW

07 **macOS version requirement raised for DPI** BREAKING

45

DPI now requires macOS 26 or later, and the `nodpi` build requires macOS 15 or later; installations on older macOS versions lose DPI or full functionality after upgrading.

— States the version cutoff but gives no migration step. v0.15.2

08 **Structured CPE matching for CVE findings** IMPROVED

45

SQLite schema 2 now evaluates structured CPE product/version applicability for CVE matching, so text-only candidates no longer produce a vulnerability finding.

— Mechanism and before/after given, no new command. v0.14.0

09 **Built-in protocol signatures replace Nmap matching** IMPROVED

40

Built-in protocol signatures now replace the Nmap database reader for service identification, retiring the Nmap probe dependency.

— Describes the replacement but no config or command detail. v0.14.0

10 **Self-contained Linux archives** IMPROVED

25

Linux archives now bundle the ELF loader and runtime libraries alongside `net` for self-contained deployment.

— Brief statement with no further operational detail. v0.14.0

└──▷ **BREAKING ON UPGRADE**

! DPI requires macOS 26+; the `nodpi` build requires macOS 15+. Installations on older macOS versions will lose DPI or full functionality after upgrade.

! Legacy v1/v2 database distribution endpoints are retired and now return HTTP 410; installations using those endpoints must switch to `net util -update-dbs`.

! Schema-1 database readers must upgrade before consuming the new layout-3 database.

! Legacy CMS and DHCP OS catalogues are disabled; HASSH catalogue enrichment is unavailable in this database snapshot.

! File capture no longer starts the Web UI by default when `--read` is set without `--http`; callers that relied on the automatic Web UI must now pass `--http` explicitly.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.48.0 NOTES ↗ CODE ↗ RANK

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

> Prowler 5.48.0 adds real-time AWS detection and remediation with automated CloudTrail-driven policy enforcement, a wave of new checks across AWS, Cloudflare, GitHub, and Vercel, and several Lighthouse AI and UI improvements including a consolidated issue side panel and early-findings streaming.

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

01 Real-time AWS detection and remediation in Prowler Cloud NEW

85

Prowler Cloud ingests CloudTrail events via EventBridge/SQS or a `LookupEvents` fallback, applying per-account policy modes `monitor`, `pr`, and `auto` (including dry-run support), with a new `/realtime` dashboard and a Manage Realtime role permission to control who can change the real-time policy. Lighthouse AI and the MCP Server gain tools to manage the event feed, read detections and actions, and revert them.

— Mechanism, policy modes and dashboard path given, no CLI command 5.48.0

02 New security checks for AWS, Cloudflare, GitHub, and Vercel

NEW

80

AWS gains a `bedrockagentcore` service with `bedrockagentcore_full_access_policy_attached`, `bedrockagentcore_payments_process_payment_role_separation`, `bedrockagentcore_runtime_vpc_configured`, plus `cloudtrail_agentcore_data_events_enabled` for AgentCore data event coverage. Cloudflare gains 13 checks across a new `account` service (`account_two_factor_enforced`, `account_member_two_factor_enabled`, `account_super_admins_limited`, `account_api_token_expiration_set`, `account_api_token_ip_restriction_enabled`, `account_api_token_recently_used`), a new `r2` service (`r2_bucket_public_dev_url_disabled`, `r2_bucket_cors_no_wildcard_origin`, `r2_bucket_custom_domain_min_tls_secure`), and a new `workers` service (`workers_script_observability_enabled`, `workers_script_workers_dev_disabled`, `workers_script_preview_urls_disabled`,

`workers_script_no_secrets_in_plain_text_bindings`), plus `zone_firewall_no_skip_all_traffic_rules`. GitHub gains `repository_secret_scanning_push_protection_enabled` to flag repositories that accept pushes containing secrets. Vercel gains `project_environment_no_secrets_with_public_prefix` (flagging secret-like environment variables exposed via prefixes such as `NEXT_PUBLIC_` or `VITE_`) and `project_linked_to_git_repository` (flagging projects that can only receive unreviewed CLI deployments).

— Lists all check IDs but no CLI invocation shown `5.48.0`

03 **Suppress Prowler Cloud promotion in CLI dashboard** NEW

75

Setting `PROWLER_NO_PROMOTIONS=1` hides the Prowler Cloud hand-off note in the OSS CLI dashboard; Prowler Cloud and Private Cloud reports never show it.

Run the local Prowler dashboard without the Prowler Cloud promotion notice, useful in automated or shared terminal environments.

```
$ PROWLER_NO_PROMOTIONS=1 prowler dashboard
```

— Exact env var and runnable command given `5.48.0`

THINNER COVERAGE BELOW

04 **Contextual Fix skill in Lighthouse AI issue panel** NEW

55

A Contextual Fix Lighthouse AI skill on the issue side panel reads open attack-path legs and findings and leads with the single change that closes the issue fastest.

— Describes behavior, lacks exact navigation steps `5.48.0`

05 **Early findings visible during running scans** NEW 55

Findings from a running first scan appear on the Findings page as they arrive, with a View Early Findings action on executing scans and critical/high findings stored first.

— Names UI action but no config `5.48.0`

06 **Issue side panel replaces separate issue page** BREAKING

55

The issue side panel now holds everything the full issue page had — Findings and Resources tabs, leg evidence, compliance frameworks, and Send to Jira — and the separate issue page is removed.

— Describes UI change and removal, no migration note `5.48.0`

07 **New compliance framework mappings for Azure and Cloudflare**

NEW 50

A GDPR framework for Azure maps Articles 25, 30, and 32 to existing Azure checks, and new Cloudflare mappings cover NCSC Cyber Essentials 3.3.

— Names frameworks and articles but no usage steps 5.48.0

08 **Demo sample data before first scan completes** NEW 50

Prowler Cloud shows labelled sample data from a demo AWS account on Overview, Findings, Resources, Compliance, Issues, and Attack Paths until the first scan completes.

— Names affected pages, no interactive steps 5.48.0

09 **Prowler Private Cloud on AWS Marketplace** NEW 50

Prowler Private Cloud is now available on AWS Marketplace as Pay as You Go and 10- and 25-account bundles, installed in the customer's own AWS account.

— Names pricing tiers, no setup steps 5.48.0

10 **Bedrock API key option for Lighthouse AI** IMPROVED 40

Prowler Cloud's Lighthouse AI settings now accept a Bedrock API key as an alternative to AWS access keys.

— Names setting but no config path or field 5.48.0

WAS THIS USEFUL?  

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ

Choose which notifications Huntress sends to ServiceNow NOTES ↗ RANK

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

> Huntress shipped finer-grained control over ServiceNow notification routing and extended SAML SSO configuration to the reseller level.

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

01 **SAML SSO at the reseller level** NEW 58

Configure SAML single sign-on at the reseller level — not just the account level — and enforce it after a successful login verifies the configuration, via Settings > Access Management, with an Enforce SSO option.

Set up and enforce SAML SSO across all accounts under a reseller, rather than configuring it account by account.

> 📍 In the Huntress portal, go to Settings › Access Management, complete the SAML SSO configuration, verify it with a successful login, then enable Enforce SSO.

— Exact UI navigation path given, mechanism explained Now Available: SSO for Resellers

02 **Notification-type filtering for ServiceNow integration** 57

BREAKING

Filter which Huntress notification types flow into ServiceNow by selecting any combination of Incidents, Escalations, Platform Actions, and Account Notices, and each ServiceNow record now shows its notification type for routing within ServiceNow workflows. Notifications without a severity now arrive in ServiceNow as Low instead of Critical (previously Critical).

— Names notification types and severity change, no config path given

Choose which notifications Huntress sen…

└──▷ BREAKING ON UPGRADE

! Huntress notifications without a severity that previously arrived in ServiceNow as Critical will now arrive as Low.

WAS THIS USEFUL?  

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.261106.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI's 7.261106.0 release adds configurable sighting de-duplication, structured custom fields and templates for Cases and Incidents, form-based knowledge submission for View Only users, and a GraphQL hardening switch, while retiring the `MERGE_USERS` feature flag.

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

01 **GraphQL secure defaults configuration** NEW 60

Add global switches for GraphQL secure defaults via the feat(graphql) configuration to harden API behaviour without per-query changes.

– Names config surface but omits exact keys or values 7.261106.0

THINNER COVERAGE BELOW

02 **Form Intake submission for View Only users** NEW 55

Users with View Only access can now submit knowledge via the Create knowledge via Form Intake capability, controlled by a dedicated permission.

– Names exact capability and notes it is permission-gated 7.261106.0

03 **MERGE_USERS feature flag removed** BREAKING 55

OpenCTI now removes the `MERGE_USERS` feature flag, making user-merge behaviour available to all deployments without explicit flag enablement.

– Names exact flag and resulting behaviour change 7.261106.0

04 **Custom de-duplication window for Sightings** NEW 33

OpenCTI now supports a custom de-duplication window for Sightings, allowing operators to tune how long the platform suppresses duplicate sighting events.

– Describes the control but no config key or default given 7.261106.0

05 **Custom fields and templates in Cases and Incidents** NEW 31

Cases and Incidents now support custom fields and templates, letting analysts capture structured, organisation-specific data inside case management workflows.

– Feature described in prose only, no field names or UI path 7.261106.0

06 **XTM Hub blog posts in news feed** NEW 25

OpenCTI now displays blog post items in the news feed from XTM Hub.

— Bare description, no mechanism or location detail 7.261106.0

WAS THIS USEFUL? 👍 👎

◆ SIEM

Wazuh

Feedback ⓘ CVE-2026-112341 CVE ↗ STABLE

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

`CVE-2026-112341`: SQL Injection in `WazuhDBBackend._substitute_params` — placeholder cross-contamination leaks agent secrets via blind oracle (CVSS 6.5, affects `>= 4.0.0`)

▾ **6 more releases in this issue** · 2026-11-06

CVE-2026-112340 CVE ↗ STABLE

`CVE-2026-112340`: wazuh-dashboard-plugins: authenticated user can override Manager API `Authorization` header via `/api/request` (CVSS 8.2, affects `>= 4.9.0`)

CVE-2026-112339 CVE ↗ STABLE

`CVE-2026-112339`: Missing TLS certificate hostname verification in the `fluent-forward` module lets any certificate issued by the configured `<ca_file>` impersonate the Fluentd collector (CVSS 5.9, affects `>= 3.9.0`)

CVE-2026-112338 CVE ↗ STABLE

`CVE-2026-112338`: WZ-2026-05 — GitHub Link next-page SSRF (CVSS 5.9, affects `>= 4.3.0`)

CVE-2026-112336 CVE ↗ STABLE

`CVE-2026-112336`: Repeated static JSON fields can exhaust memory and terminate wazuh-analysisd (affects `>= 3.0.0`)

CVE-2026-112335 CVE ↗ STABLE

▶ Watch video

`CVE-2026-112335`: Wazuh Windows Agent - Stack-Based Buffer Overflow in read_snortfull() Leads to SYSTEM-Level Code Execution Primitive (CVSS 5.5, affects `>= 1.0.0`)

CVE-2026-112334 CVE ↗ STABLE

`CVE-2026-112334`: Out-of-bounds write in wazuh-execd when loading more Active Response commands than the table can hold (CVSS 4.9, affects `>= 1.0.0`)

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

Install drivers for JDBC Unity Catalog connections from Maven (2026-11-06)

BLOG ↗ RANK

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

> Databricks shipped three distinct capabilities this window: Maven-based JDBC driver installation for Unity Catalog connections, a new Beta for collaborative Unity Catalog wikis, and general availability for metric view listings in Databricks Marketplace.

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

01 **Unity Catalog wikis in Catalog Explorer** NEW 70

Databricks now supports wikis on Unity Catalog tables, letting data owners and stewards draft and publish documentation manually or with Genie Code, while viewers can comment on, suggest edits to, and react to published wikis in Catalog Explorer. Account admins can control access to this Beta feature from the Previews page (Account Console > Settings > Previews > Unity Catalog wikis).

Enable the Unity Catalog wikis beta for your account so data stewards can begin documenting tables.

> 📍 In the Databricks Account Console, go to Settings › Previews, find Unity Catalog wikis , and enable it.

— Names mechanism, roles, and exact enablement path

Unity Catalog wikis are in Beta (2026-1…

02 **Maven driver install for JDBC Unity Catalog connections** NEW 65

JDBC drivers for Unity Catalog connections can now be installed from a Maven coordinate instead of being uploaded to a volume. This requires serverless compute, Databricks Runtime 19 or above, or SQL warehouses 2026.43 or above.

— Clear requirements but no exact UI/CLI steps given

Install drivers for JDBC Unity Catalog …

THINNER COVERAGE BELOW

03 **Metric view listings GA in Marketplace** IMPROVED 45

Metric view listings are now generally available in Databricks Marketplace, letting providers publish governed business metrics by adding metric views to a listing. This is

available to Databricks recipients only.

— Describes what changed but no steps or scope details

Metric view listings are generally avai…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Endpoint Management (UEM/RMM)

N-able N-central Feedback ⓘ

Patch Management Engine (PME) 2.14.3 General Availability (2026-11-06) BLOG ↗

N-able n-central is a secure endpoints and automate IT operations at scale. N-central delivers unified endpoint management with advanced RMM capabilities for complex environments.

NEW SUPPORTED PLATFORM Windows 11 26H2 Feature Update support via Enablement Package

WAS THIS USEFUL? 👍 👎

ITarian

Feedback ⓘ

Release Notes: ITarian Platform Update - November 11, 2026 (2026-11-06) BLOG ↗

RANK

Streamline IT operations with ITarian's all-in-one IT management software. Automate tasks, enhance security, and boost efficiency.

> ITarian's November 2026 update centers on a more configurable Service Desk dashboard with date-range and department filtering plus chart exports, alongside Excel-style ticket filtering, device list customization, and a per-device chat view.

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

01 **Customizable Service Desk dashboard with filters and charts**

NEW 78

Filter the Service Desk dashboard by date range (Today , Yesterday , This Week , Last Week , Last 7 Days , Last 30 Days , This Month , Last Month , Last 365 Days , or Custom range) and by Department, with results applied together across stat cards, chart, and the Statistics table. Switch the chart between Line, Bar, and Pie views and export the dashboard to CSV or PDF.

— Names every filter option and chart/export format

Release Notes: ITarian Platform Update …

02 **Ticket list filtering and display enhancements** NEW

65

Ticket columns gained Excel-style per-column filters — select one or more values, combine filters across columns, and clear individual columns or all filters at once via Clear all . The ticket inbox can also be filtered by organization, and the top menu bar and Ticket List column headers can now be pinned while scrolling across Inbox, My Tickets, Unassigned, Overdue, and Closed views.

— Describes mechanism but no exact UI path given Release Notes: ITarian Platform Update …

THINNER COVERAGE BELOW

03 **Custom column order in Device List views** NEW 50

Set the order of columns in a Device List custom view, saved with the template and shared across all admins on the portal.

— Explains behaviour and persistence but no config name

Release Notes: ITarian Platform Update …

04 **Device-specific chat conversations** NEW 44

Open chat from a specific device to see only that device's conversation instead of a global inbox spanning all users and devices.

— Clear behaviour change, lacks exact navigation path

Release Notes: ITarian Platform Update …

05 **Service Desk data export to CSV** NEW 28

Export Service Desk data to CSV for analysis and reporting.

— Bare statement with no scope or location detail

Release Notes: ITarian Platform Update …

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ What's New in Iru | November 6th NOTES ↗

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

Iru adds Remote Shell to EDR for all customers and rolls Auto Apps update notifications into a single daily digest.

Iru EDR now includes Remote Shell, available to all customers.

Auto Apps update notifications now consolidate into a single daily digest instead of individual alerts.

NEW SUPPORTED PLATFORM Korean language support

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

iDenfy

Feedback ⓘ updated · Nov 7, 2026 NOTES ↗ RANK

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

> iDenfy's SDKs gained standalone bank card verification and new Brazilian document checks, alongside smaller Android and iOS verification UX additions.

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

01 **Bank Card Verification standalone mode** NEW 55

Bank Card Verification standalone mode is now available on Android `9.2.4` and iOS `9.2.9`, allowing card verification to run independently of the full identity verification flow.

— Names mechanism and exact SDK versions but no integration steps. snapshot-20261107

02 **CPF and CNH digital ID checks in SDKs** NEW 41

CPF check support was added to Android `9.2.4` and iOS `9.2.9`, and CNH digital ID support was added to Android `9.1.1` and iOS `9.1.1`, extending Brazilian document and ID checks available to the mobile SDKs.

— Names both checks and SDK versions but no usage detail. snapshot-20261107

03 **Country mismatch dialog in Android SDK** NEW 33

Android `9.2.4` adds a new country mismatch dialog to surface document-country discrepancies during verification.

— States purpose but no further mechanism or config. snapshot-20261107

04 **iPhone Duo support in iOS SDK** NEW 16

iOS `9.2.9` adds iPhone Duo support.

— Bare one-line mention with no further explanation. snapshot-20261107

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ

Application Logs now group sign-ins and sign-ups (2026-11-06) 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 shipped grouped log entries and new email-event visibility in Application Logs, one-time authenticated browser links for AI agents, and custom per-user/device session expiry controls.

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

01 Grouped sign-in/sign-up events in Application Logs BREAKING

65

Clerk's Application Logs now group sign-ins and sign-ups into a single row per event, and `GET /v1/logs` returns these groups by default instead of individual log lines. Callers that expected individual log lines per event will receive a different response shape without any documented opt-out.

— Endpoint named but no opt-out flag or migration path given

Application Logs now group sign-ins and…

02 Email events in Application Logs and Backend API NEW

65

Email events can now be queried alongside every other event type via `GET /v1/logs` in the Backend API, and viewed next to other events on the Application Logs page in the Dashboard.

Retrieve recent email events from your application logs to audit delivery or debug email flows.

```
$ curl -X GET 'https://api.clerk.com/v1/logs' \
  -H 'Authorization: Bearer <secret_key>'
```

— Runnable curl example against the named endpoint

Email events in Application Logs and th…

THINNER COVERAGE BELOW

03 One-time authenticated browser links for agents NEW

45

Clerk adds agent browser sessions so an MCP server can issue a one-time sign-in link that opens a browser already authenticated as the user, letting agents perform browser-based tasks on the user's behalf.

– No concrete endpoint or command shown for issuing the link

Browser sessions for agents (2026-11-06)

---

04 **Custom session limits per user or device** NEW 45

Clerk adds per-user and per-device custom session expiry, settable for individual sessions via a new Backend API endpoint, with session limits also controllable through session token shortcodes.

– Endpoint path and shortcode names not specified

Set custom session limits per user or d…

└──▷ **BREAKING ON UPGRADE**

! The `GET /v1/logs` endpoint now returns grouped log entries by default; callers that expected individual log lines per event will receive a different response shape without opting out of grouping.

WAS THIS USEFUL?  

◆ Secure Access (VPN / SASE / ZTNA)

amneziawg-installer

Feedback ⓘ v6.2.1 NOTES ↗ RANK

1 VPN server on a fresh Ubuntu or Debian VPS with one command. Made for networks that block WireGuard.

> amneziawg-installer added a security mitigation that disables apt download retries during installation to close a package-substitution vulnerability.

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

01 **Apt retry mitigation for CVE-2026-109518** NEW 88

The installer now writes `/etc/apt/apt.conf.d/99zz-amneziawg-installer-noretries` at the start of every run to disable apt download retries and mitigate CVE-2026-109518, deleting the file on exit; `--uninstall` also removes it if a run was interrupted. It also warns when apt retries remain enabled via another apt setting despite this file being in place.

Upgrade apt itself without retries on an already-installed server, to close the same CVE-2026-109518 window the installer now guards against.

```
$ apt-get -o Acquire::Retries=0 update && apt-get -o
Acquire::Retries=0 install --only-upgrade apt
```

— Names exact config file path and uninstall behaviour, but no flag to control it v6.2.1

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

warden

Feedback ⓘ v0.22.0 NOTES ↗ RANK

The secure gateway connecting AI agents to enterprise systems.

> Warden v0.22.0's headline is agent identity governance — entities, aliases, sponsors and an `unknown_aliases` enforcement mode — alongside an OIDC login flow, a Microsoft Graph provider, certificate/KMS-backed Azure and OAuth2 authentication, and cheaper policy compilation, shipped together with a large set of breaking changes to the role, token, discovery, MCP-resources, and naming model.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 Agent identity governance with entities and sponsors NEW

95

Manage agent identity with `warden identity entity` and `warden identity alias`, and surface ungoverned agents via `warden identity entity list -ungoverned`. The `unknown_aliases` setting controls unclaimed agents (`observe` logs why `deny` would refuse each request, `deny` actively refuses unclaimed and ungoverned agents), and `governed_only` roles are restricted to what a claimed agent's entity groups explicitly allow. Agent sponsors — virtual groups of humans who accept an agent — are managed with `warden identity sponsor`, without requiring a Warden role.

Find agents that have no sponsor or entity so you can decide whether to claim or block them before switching `unknown_aliases` to `deny`.

```
$ warden identity entity list -ungoverned
```

— Full named command set and mechanism, ready to run v0.22.0

02 Auth mount roles replaced by top-level role API BREAKING

90

`auth/<mount>/role/*` now returns `410`; roles must instead be created with `warden role create` at `sys/roles/<name>` and granted by groups per request. Mount keys `default_role`, `groups_claim`, `group_policy_prefix`, `metadata_claims`, `metadata_mappings`, `token_type`, and the provider's `user_auth_role` are refused. Agent `jwt` and `spiffe` mounts now require `bound_audiences`, and `max_age` moves from the role to the mount. In mixed-version clusters, a v0.21.0 node cannot read PS256 capabilities, roles, or groups.

— Exact endpoint, replacement command and refused keys named v0.22.0

03 **Human login via OIDC and built-in operator roles** NEW

81

Sign people in with `warden login` through the new `oidc` auth method, and manage sessions with `warden logout`, `warden role use`, and `warden print token`. Four built-in operator roles (`warden-admin`, `warden-provider-admin`, `warden-security-admin`, `warden-viewer`) can be enabled per namespace using the `enable_builtin_roles` setting.

Sign in as a human operator through your IdP and switch to the read-only built-in role for your namespace.

```
$ warden login
warden role use warden-viewer
```

— Named commands and roles, missing permission detail v0.22.0

04 **Memory caps for token and credential caches** NEW

81

Set `token_cache_megabytes` and `credential_cache_megabytes` to cap how much memory Warden's token and credential caches use. At the default 100 MiB limit, the token cache holds approximately 113,000 service tokens.

Cap Warden's token and credential caches to specific memory limits so the process stays within your container's memory budget.

```hcl
token_cache_megabytes = 200
credential_cache_megabytes = 150
```

— Exact config keys, default limit, and runnable example v0.22.0

05 **Certificate and KMS-backed auth for Azure and OAuth2 sources**

IMPROVED

74

Azure sources can now authenticate with a certificate whose private key stays in a KMS. OAuth2 sources can authenticate with `client_secret_basic`, `private_key_jwt`, or `kms_private_key_jwt`, and an issuer audience can now be set on `oauth2` and `token_exchange` sources.

— Names each method but omits the config keys to set them v0.22.0

06 **New token types invalidate prior sessions** BREAKING

68

Token types are now `identity_token` and `service_token` (prefix `wds.`); tokens issued by a v0.21.0 auth mount are refused, requiring agents to log in again. Root and other `cws.` tokens from no-longer-present mounts will not resolve until they expire.

— Names token prefixes; migration path is only re-login v0.22.0

07 **Microsoft Graph provider and Bedrock Runtime gateway routing**

NEW 65

The new `msgraph` provider reaches Microsoft 365 as the calling user via the Entra on-behalf-of flow. Bedrock Runtime model calls are now routed through the `aws` gateway, bounded by the mount's `timeout`.

— Names provider and flow but no setup steps given v0.22.0

08 **Discovery now requires credentials and caps batches at 50**

62

BREAKING

Discovery now requires a credential, returning `401` without one, and is capped at batches of 50; `503` is returned when no mount admitted the credential.

— Status codes and cap named, no remediation example v0.22.0

THINNER COVERAGE BELOW

09 **MCP `resources` family gains skills and directory methods**

59

BREAKING

The MCP `resources` family now covers `skills/get` and `resources/directory/read`; a contract allowing every method with no `resources` block now refuses those methods.

— Names methods but not how to update contracts v0.22.0

10 **hvault credential source requires auth_method, blocks wid: principals**

59

BREAKING

`hvault` credential source edits now require `auth_method`; `wid:` user principals are refused; impersonation and routing headers are stripped.

— Names the required field and refused principal type v0.22.0

11 **GCP and Azure forwarding restricted to allowed host suffixes**

57

BREAKING

`gcp` and `azure` sources now forward only to `allowed_host_suffixes`, and `azure` refuses Microsoft Graph `$batch` requests.

— Names the config key and the refused request type v0.22.0

12 **Audit logging floor removes default header and body capture** 56

BREAKING

Audit now has a floor: no headers or bodies are logged by default; `salt_fields` extends default salting; unknown paths are refused at enable time.

— Names the config key but not full salting scope v0.22.0

13 **Reserved and canonical name rules for policies and roles** 52

BREAKING

Policy names `default` and `default-*`, and role names `warden-*`, are now reserved and may not be used. Non-canonical policy names and percent-encoded dot or empty path segments now return `400`.

— Names reserved patterns and error code, no remediation v0.22.0

14 **Built-in discovery and sponsor policies** NEW 42

Built-in policies `default-cbp`, `default-mcp`, and `default-sponsor` are now available for discovery and sponsor workflows.

— Names the policies but not what they enforce v0.22.0

15 **Nested namespaces written through an ancestor were misplaced** 40

BREAKING

Nested namespaces written through an ancestor were misplaced and must be recreated manually; no automated repair is provided.

— States the bug and fix but no reproduction steps v0.22.0

16 **`request.data_plane` field in CEL policies** NEW 35

CEL policy expressions can now use `request.data_plane` to distinguish data-plane requests from other request types.

— Names the field but not the distinguishing behaviour v0.22.0

17 **Policy compilation 2.7x cheaper** IMPROVED 35

Policy compilation is now about 2.7 times cheaper, enabling larger policy sets without a proportional latency cost.

— A number with no mechanism or benchmark detail v0.22.0

└──▷ BREAKING ON UPGRADE

- ! `auth/<mount>/role/*` now returns `410`; roles must be created with `warden role create` at `sys/roles/<name>` and granted by groups per request. Mount keys `default_role`, `groups_claim`, `group_policy_prefix`, `metadata_claims`, `metadata_mappings`, `token_type`, and the provider's `user_auth_role` are refused.
- ! Agent `jwt` and `spiffe` mounts now require `bound_audiences`; `max_age` moves from the role to the mount.
- ! Token types are now `identity_token` and `service_token` (prefix `wds.`); tokens issued by a v0.21.0 auth mount are refused, requiring agents to log in again. Root and other `cws.` tokens from no-longer-present mounts will not resolve until they expire.
- ! Policy names `default` and `default-*`, and role names `warden-*`, are reserved and may not be used.
- ! Discovery now requires a credential (`401` without one) and is capped at batches of 50; `503` is returned when no mount admitted the credential.
- ! The MCP `resources` family now covers `skills/get` and `resources/directory/read`; a contract allowing every method with no `resources` block now refuses those methods.
- ! `gcp` and `azure` forward only to `allowed_host_suffixes`; `azure` refuses Graph `$batch` requests.
- ! Non-canonical policy names and percent-encoded dot or empty path segments return `400`.
- ! `hvault` credential source edits now require `auth_method`; `wid:` user principals are refused; impersonation and routing headers are stripped.
- ! Audit now has a floor: no headers or bodies are logged by default; `salt_fields` extends default salting; unknown paths are refused at enable time.
- ! Nested namespaces written through an ancestor were misplaced — they must be recreated; no automated repair is provided.
- ! In mixed-version clusters, a v0.21.0 node cannot read PS256 capabilities, roles, or groups; the token cache holds approximately 113,000 service tokens at its default 100 MiB limit.

WAS THIS USEFUL?  

◆ Backup & Recovery

Acronis Cyber Infrastructure

Feedback ⓘ 8.0.0-401

NOTES ↗

Acronis Cyber Infrastructure is hyperconverged software that pools storage, compute, and networking on commodity servers, often as a backup storage target.

Acronis Cyber Infrastructure 8.0 adds single-node deployment, a guided setup wizard, VMware migration GA, API access controls, compute performance profiles, and significant storage throughput gains.

Adds a single-node deployment model: a single-node cluster with local storage, each with its own admin panel, self-service panel, and alerting capabilities.

The VMware migration service is now generally available for administrators and self-service users, with support for project-level access control, CPU and RAM hotplug, and migration of VMs with multiple volumes.

Adds a guided initial setup wizard that lets administrators configure compute and storage services through a single guided workflow on first login.

Adds an option to control API access through user permissions, allowing administrators to specify which users can access the API and require API credentials for authentication.

Adds support for selecting a performance profile for the compute cluster, defining how much CPU and RAM compute services reserve on each node.

+4 more

NEW SUPPORTED PLATFORM Ubuntu 26.04 LTS guest OS support

NEW SUPPORTED PLATFORM SUSE Linux Enterprise Server (SLES) 16 guest OS support

RETIRED SUPPORTED PLATFORM CentOS 6 guest OS support deprecated

NEW INTEGRATION Official Huawei storage driver 25.3.0

UPDATED INTEGRATION Hitachi Cinder driver updated to version 2.5.0

UPDATED SUPPORTED PLATFORM Linux kernel updated to version 6.12

UPDATED SUPPORTED PLATFORM Kubernetes updated to version 1.36.2

UPDATED SUPPORTED PLATFORM Fedora CoreOS base image for managed Kubernetes updated to version 44

UPDATED SUPPORTED PLATFORM NVIDIA drivers updated to R595 branch; vGPU software updated to version 20

└──▷ BREAKING ON UPGRADE

- ! Version 8.0 introduces a new OS and kernel stack based on RHEL 10; verify cluster hardware meets compatibility requirements before upgrading.
- ! Nodes with NVIDIA GPUs or vGPUs must have NVIDIA vGPU software driver version 595.58.03 or later installed before upgrading to version 8.0.
- ! Storages attached with the legacy Huawei driver must be manually switched to the new official Huawei driver version 25.3.0 after upgrading.
- ! CentOS 6 guest OS support is deprecated in this release.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

h5i

Feedback ⓘ v0.5.4 NOTES ↗ RANK

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

> h5i's first public release ships a plugin-based toolkit for web app red-teaming: Chromium-backed browser automation with full request auditing, an HTTP replay workbench, an endpoint-discovery recon plugin, OpenAPI-coverage regression testing, a session dashboard, and a sandboxed execution environment.

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

01 **Security regression testing against OpenAPI coverage** NEW

90

Run `h5i test .h5i-tests/tests --target <url> --openapi <spec> --min-coverage <pct>` to audit a target and enforce a minimum percentage of OpenAPI endpoints exercised before CI passes; supply `--in` to the test plugin to add an out-of-browser witness to the request log on a sandboxed host.

Run a full audit of a staging target and enforce that at least 80% of OpenAPI endpoints are exercised before CI passes.

```
$ h5i test .h5i-tests/tests --target https://staging.example.com
--openapi openapi.yaml --min-coverage 80
```

— Runnable command with concrete flags and example values v0.5.4

02 **Browser automation and request logging** NEW 85

`h5i browser open`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` capture, policy-check, and log every browser request in a Chromium session driven through agent-browser, so you can inspect every request made (and what was refused) and dump the full session audit for review.

Inspect every request a Chromium session made and what was refused, then dump the full session audit for review.

```
$ h5i browser open https://example.com
h5i browser requests
h5i browser audit
```

— Named commands and mechanism, runnable example given v0.5.4

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

Install with `h5i plugin install recon --from ./h5i-recon` to record endpoints, parameters, and candidate URLs a target exposes, with discoveries confirmed only when a live request answers; crawl a target with `h5i recon crawl --max-requests 200`.

Install the recon plugin and crawl a target to discover exposed endpoints and candidate URLs.

```
$ h5i plugin install recon --from ./h5i-recon
h5i recon crawl --max-requests 200
```

— Named plugin, flag, and runnable crawl command v0.5.4

04 **Websec workbench plugin for HTTP replay** NEW 60

Install with `h5i plugin install websec --from ./h5i-websec` to replay, modify, and compare HTTP requests, enforced through the same policy and budget controls used by the browser automation.

— Install command named but no usage example shown v0.5.4

THINNER COVERAGE BELOW

05 **Sandboxed environment for isolated testing** NEW 50

Start a sandboxed environment with filesystem and network restrictions using `h5i box --name <name> --profile browser`.

— Command and flags named but scope of restrictions unexplained v0.5.4

06 **Local dashboard for session monitoring** NEW 35

Run a local dashboard to monitor browser sessions and agent activity with `h5i ui`.

— Single command named, no further mechanism given v0.5.4

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Microsoft Sysinternals

Feedback ⓘ

Autoruns for Linux, ProcExp for macOS, macOS tools, Linux tools, ProExp, ProcMon, PsTools

BLOG ↗ RANK

Microsoft sysinternals is a our community is a collaborative place to learn about our products and ask questions! Check out our most popular conversations!

> Sysinternals expanded cross-platform coverage with new Linux and macOS ports of Autoruns and Process Explorer, while also speeding up Process Explorer, adding dictation support to ZoomIt for macOS, and switching RDCMan to Credentials UI authentication.

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

01 **Autoruns for Linux v1.0** NEW 45

A new Linux port of Autoruns (v1.0) lets users view and manage boot and login auto-start configuration on Linux systems.

– Names version and platform but no usage detail Autoruns for Linux, ProcExp for macOS, …

02 **Process Explorer for macOS v1.0** NEW 43

A new macOS port of Process Explorer (v1.0) brings the process diagnostic and management utility to macOS.

– Names version and platform but no usage detail Autoruns for Linux, ProcExp for macOS, …

03 **Dictation for ZoomIt screen snip captions** NEW 43

ZoomIt 12.3.2 for macOS now supports dictation for captions on screen snips.

– Names version and feature but no usage steps Autoruns for Linux, ProcExp for macOS, …

04 **RDCMan Credentials UI authentication** IMPROVED 35

RDCMan v3.23 now uses Credentials UI for authentication.

– Names version and change but no further mechanism

Autoruns for Linux, ProcExp for macOS, …

05 **Faster Process Explorer startup** IMPROVED 25

Process Explorer v17.21 now starts faster.

– Bare improvement claim with no mechanism or numbers

Autoruns for Linux, ProcExp for macOS, …

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 7, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue