THE CYBER TOOLCHAIN

NO. 136

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 11, 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 221 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 63 of 64 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

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. OpenSandbox's Credential Vault Exchange resolves secrets per request across all five SDKs, so

vim ? help

sandboxes stop holding long-lived keys. Manifest now traces which direct dependency pulled in a vulnerable transitive component. Hayabusa scans EWF and VHDX images natively, skipping the mount-and-extract step before a timeline.

> PROTECT
>
> **Stop handing sandboxes long-lived credentials**
>
> OpenSandbox's Credential Vault Exchange resolves credentials dynamically at request time across all five SDKs rather than baking them into the sandbox environment, and Tenuo adds opt-in consume-on-use approvals so an agent's authorization dies after one call. Chainguard's SCIM 2.0 provisioning closes the other end — deprovisioning follows the IdP instead of waiting on a manual ticket.
>
> OpenSandbox · tenuo · Chainguard

> PROTECT
>
> **Find out which direct dependency dragged in the bad package**
>
> Manifest's Components tab now maps transitive dependencies back to the direct one that introduced them, which is the step that turns an unactionable CVE on a package nobody chose into a specific version bump. JFrog's Threat Hub does the inverse from the intel side, resolving a vulnerability to the artifacts across the org that carry it; Skylos extends supply-chain scanning to malware feeds and install hooks, where the compromise arrives before any CVE exists.
>
> Manifest · JFrog · skylos

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Wireless & RF

wifite

Feedback ⓘ 3.0.0 NOTES ↗

Wifite is an automated wireless network auditing tool that attacks multiple WEP, WPA, and WPS encrypted networks simultaneously using various cracking techniques.

wifite 3.0.0 adds MAC randomization flags, a per-run capture and crack cache, Scapy PMKID fallback, scan diagnostics, and a radio-state supervisor that owns every interface mutation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.0.0 https://github.com/kimocoder/wifite2.git
# already have the repo? check out this version:
$ git checkout 3.0.0
```

Randomize the interface MAC address with `--random-mac` for a fully random address or `--random-mac-vendor` to randomize only the device part while keeping vendor bytes.

Skip already-captured networks in a run with `--ignore-captured` to avoid redundant work.

Point `--dict` at a directory of wordlists instead of a single file.

Check the installed version with `--version`.

A run-scoped capture cache lets a PMKID run publish its pcapng so the WPA stage can reuse it rather than capturing again.

+8 more

NEW SUPPORTED PLATFORM 6 GHz Wi-Fi band support

NEW SUPPORTED PLATFORM Broadcom FullMAC / Nexmon adapter support

UPDATED SUPPORTED PLATFORM hcxdumptool/hcxtools 8 and hashcat 8 support

└──▷ BREAKING ON UPGRADE

! hcxdumptool and hcxpcapngtool 7.0.0 are now required minimums; older versions will no longer work.

! WPA-FT now requires hashcat 8.0.0 or newer.

! hcxpcapngtool `--bssid` and `--essid` flags are no longer passed; filtering is now done in Python, so capture pipelines that relied on those flags will behave differently.

! Automatic self-update from `origin/master` is removed; running the old update flow will no longer apply changes.

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Strix

Feedback ⓘ v1.7.1 NOTES ↗

Strix adds timing telemetry via `STRIX_RUN_RECORD_TIMINGS`, parallel tool calls via `STRIX_PARALLEL_TOOL_CALLS`, host device passthrough via `STRIX_DOCKER_SANDBOX_DEVICES`, and arbitrary LLM request body injection via `LLM_EXTRA_BODY`; the TUI gains a targets sidebar and sets the terminal title; reachability now has three states; image-unaware models are auto-detected and images scrubbed; dependency reports accept `code_locations`; and sub-agents have soft per-skill turn limits.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.7.1 https://github.com/usestrix/strix.git
# already have the repo? check out this version:
$ git checkout v1.7.1
```

Record every LLM request and tool call's timing into `run.json` by setting the `STRIX_RUN_RECORD_TIMINGS` environment variable.

Enable opt-in parallel tool calls with the `STRIX_PARALLEL_TOOL_CALLS` environment variable.

Pass host devices into the sandbox container via the `STRIX_DOCKER_SANDBOX_DEVICES` environment variable.

Inject arbitrary extra fields into LLM request bodies via the `LLM_EXTRA_BODY` environment variable.

Add a `-r` flag and a title column to the resume picker to identify saved runs at a glance.

+10 more

NEW DETECTION XSS via HTML injection and normalization

NEW DETECTION Email identity canonicalization drift

NEW DETECTION Query-language and JSON parser differentials

NEW INTEGRATION Opengrep static analysis integration

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

adscan

Feedback ⓘ v14.0.0 NOTES ↗ RANK

> ADscan v14.0.0 ships a combined scan-and-report command, configurable scan policies with pacing controls, plan-file driven runs, and a no-rescan recheck command, alongside new AD CS ESC12 detection, richer endpoint-defense and attack-path reporting, and several breaking changes to `adscan ci` and `adscan execute`.

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

01 **Scan policies and pacing controls** NEW 95

`adscan ci --policy` selects a scan policy from three built-ins — `lab_ctf` (every check at Fast pace), `audit` (self-rolling attack steps at Normal pace), and `production_sensitive` (no password guessing, no directory changes, Polite pace). `adscan policy` manages custom scan policies on the host, including `policy review-new` to see which checks arrived after a policy was saved. `--pacing` sets scan speed independently across four levels — Fragile, Polite, Normal, Fast — savable as a pace profile, with the policy capping maximum speed and ADscan throttling automatically when a server struggles. The `ADSCAN_PORTSCAN_NMAP_TIMING` environment variable caps nmap timing in port scans to a maximum of `-T4`.

— Enumerates every flag, policy name, pace level and env var. v14.0.0

02 **Combined scan-and-report command** NEW 90

`adscan ci --generate-report` runs a full scan and renders the client deliverable in one step, accepting every option `adscan deliver` takes — document selection, frameworks, theme, branding, client name and logo, custom cover and back page, and outputs. A scan whose results are good but whose report fails to render exits with code `4`.

— Names exact flag, inherited options, and exit code. v14.0.0

03 **Plan-file driven scans via `adscan run`** NEW 90

`adscan run --plan plan.json` drives a scan, connection check, or deliverable kit from a reviewable file, referencing credentials by name and streaming versioned progress through `--events-file`.

Run a scan from a reviewable plan file that keeps credentials out of the command line and streams progress for a CI system to consume.

```
$ adscan run --plan plan.json --events-file events.jsonl
```

— Command, flags, and a runnable usage example given. v14.0.0

04 **Breaking changes to `adscan ci` and `adscan execute`** BREAKING

85

`--type` and `set type` on `adscan ci` no longer change scan behaviour; they now print the equivalent policy command instead, and existing workspaces migrate automatically to the built-in policy their former type mapped to. `adscan execute` moves to the PRO tier alongside `adscan ci`, so free-tier users must use `adscan start`. Both `adscan ci` and `adscan execute` now keep their workspace by default; pass `--ephemeral` to delete it on exit, reversing the previous default.

— Names flags, commands and exact migration/default behaviour. v14.0.0

05 **Workspace secret purge command** NEW

75

`workspace purge-secrets` removes recovered passwords, hashes, keys, tickets and certificates from a workspace while keeping its findings and history intact; the removal is recorded with the operator who ran it.

— Names command, scope, and audit behaviour. v14.0.0

06 **Endpoint Defense assessment reporting** IMPROVED

70

The Endpoint Defense section now reports, per host, which endpoint defense product was active, any product installed but not running, LSA protection and Credential Guard settings, and what the assessment proved — with a recommendation to deploy EDR on Tier 0 hosts that have none.

— Detailed field list but no command surface to act on. v14.0.0

07 **Recheck findings without a new scan** NEW

65

`adscan recheck` re-verifies an earlier scan's findings without running a new scan, and can repeat a single attack step or walk a single path read-only.

— Command named with clear usage, no options detailed. v14.0.0

08 **AD CS ESC12 detection** NEW

65

ADscan now detects AD CS ESC12, where a certificate authority's YubiHSM2 password sits in clear text on the CA host, identified using two read-only registry reads.

— Vulnerability and detection mechanism named, no CLI surface given. v14.0.0

09 **New exit codes for skipped and failed CI scans** NEW

65

The free tier now exits with code `42` after the PRO notice so a pipeline never reads a skipped scan as a success; `adscan ci` exits `3` when a collection or path search does not complete.

— Exact exit codes named with clear CI consequence. v14.0.0

10 **Attack-path execution display with MITRE IDs** NEW 60

Attack-path execution now shows the plan first, then each step as it starts and settles with its MITRE technique ID; a step-through mode asks before each step.

— Describes UX change, no flag named for step-through mode. v14.0.0

THINNER COVERAGE BELOW

11 **Compliance framework editions and verified-clear scorecards**

IMPROVED 45

Both reports now state which edition of each compliance framework they map against, and scorecards report verified-clear status on attack surfaces the scan enumerated and found nothing abusable on.

— Reporting change described only in prose, no named surface. v14.0.0

└──▷ **BREAKING ON UPGRADE**

! `--type` and `set type` on `adscan ci` no longer change scan behaviour; they print the equivalent policy command instead. Existing workspaces migrate to the built-in policy their former type mapped to.

! `adscan execute` moves to the PRO tier alongside `adscan ci`; free-tier users must use `adscan start`.

! `adscan ci` and `adscan execute` now keep their workspace by default; pass `--ephemeral` to delete it on exit (previously the default was to delete it).

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

nerva

Feedback ⓘ v1.84.3 NOTES ↗

└──▷ GET THIS VERSION ▾

Go

```
$ go install github.com/praetorian-inc/nerva/cmd/nerva@v1.84.3
```

git

```
$ git clone --branch v1.84.3 https://github.com/praetorian-
inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.84.3
```

Go module path matches the repo · checked 11 Nov Wrong command?

NEW DETECTION Shared remote-access service fingerprinting

WAS THIS USEFUL? 👍 👎

myip

Feedback ⓘ v8.1.0 NOTES ↗ RANK

> myip's v8.1.0 release adds a Bulk IP Lookup workflow with geolocation, filtering, and export, plus a per-IP scenario friction score covering 10 use cases, alongside smaller refinements to DNS Resolution and Network Connectivity tools.

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

01 **Bulk IP Lookup with sorting, filtering and export** NEW

55

Users can paste IP lists or raw server logs to geolocate all IPs at once, then sort, filter, star, summarize results by country and ASN, map them, and export the results.

— Describes workflow steps but no UI path or export format named v8.1.0

02 **Scenario friction score on IP cards** NEW

40

IP cards now show a Scenario friction estimate — the extra checks or restrictions an IP may face across 10 scenarios including sign-up, payments, and streaming.

— Explains concept with scenario count but no mechanism detail v8.1.0

03 **More detailed DNS Resolution results** IMPROVED

7

The DNS Resolution tool now returns more detailed results.

— Bare statement with no specifics on what changed v8.1.0

04 **More importable lists in Network Connectivity** IMPROVED

7

Network Connectivity now supports more importable lists.

— Bare statement with no named formats or lists v8.1.0

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

Hak5 Bash Bunny

Feedback ⓘ sold out · Nov 10, 2026 STORE ↗

Bash Bunny is a Linux-based USB attack platform that mimics keyboards, storage, serial and Ethernet devices to run multi-vector payloads against a computer within seconds of plugging in.

HotPlug Attack Combo Kit (bash-bunny-hotplug-combo) sold out at `shop.hak5.org` on 10 Nov 2026.

WAS THIS USEFUL? 👍 👎

Hak5 Shark Jack Display

Feedback ⓘ

sold out · Nov 10, 2026 STORE ↗

Shark Jack Display is a jack into a network and instantly run advanced recon, exfiltration, attack and automation payloads.

Incognito Black, single (sjd-incognito-single) sold out at `shop.hak5.org` on 10 Nov 2026.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ WAF, Bot & API Protection

safeline

Feedback ⓘ v9.4.4 NOTES ↗

SafeLine is a self-hosted WAF (Web Application Firewall) that protects websites from web attacks like SQL injection, XSS, and malicious traffic.

SafeLine adds ASN-based custom rule operators, a Beta upgrade channel, custom upstream health-check paths, dynamic DNS resolution for upstreams, and a console announcement banner.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v9.4.4 https://github.com/chaitin/SafeLine.git
# already have the repo? check out this version:
$ git checkout v9.4.4
```

Block attack sources by autonomous system number using new `Belongs to ASN` and `Does Not Belong to ASN` match operators in custom rules (Lite/Pro); matched logs now include ASN info.

Opt in to a Beta upgrade channel from the console to try new features early, then upgrade smoothly to stable once published.

Configure a custom check path for upstream health checks so availability is accurately detected when your app is not served from the root path.

Dynamic DNS resolution for domain-based upstreams (Pro) now updates the resolved IP automatically when DNS records change, so failover no longer requires a manual reload.

A new announcement banner at the top of the console now shows release updates and important notices, with options to view detail, dismiss temporarily, or permanently hide.

UPDATED SUPPORTED PLATFORM IP geolocation database update

WAS THIS USEFUL? 👍 👎

fcaptcha

Feedback ⓘ v1.48.0 NOTES ↗ RANK

> fcaptcha shipped a new scored detection rule for fast-start stealth-CDP sessions, enforced consistently across its Node, Go, and Python runtimes.

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

01 **Stealth-CDP session scoring rule** NEW 73

fcaptcha now assigns a production score floor of `0.8` to fast-start stealth-CDP sessions when at least 30 pointer movement samples and 30 coalescing samples are present, fewer than 5% zero movement deltas are observed, a CDP console consumer is detected, and first interaction occurs under 250 ms after page load. These stealth-CDP score boundaries are now enforced identically across the Node, Go, and Python runtimes.

— Detailed thresholds given but no config/API to invoke it v1.48.0

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ Fastly Terraform Provider (beta) 0.2.9 (2026-11-10)

BLOG ↗ RANK

Fastly is a cloud edge platform providing CDN, DDoS protection, load balancing, and real-time content delivery to accelerate and secure web applications.

> Fastly shipped IAM role-based authentication for AWS S3 backends alongside new metrics to monitor those IAM-authorized requests, added SubtleCrypto `deriveBits` / `encrypt` / `decrypt` support to the JavaScript SDK via StarlingMonkey v0.5.0, introduced dry-run and backup-write modes plus a version guardrail for the Terraform provider's migration tool, and changed Bot Management so category actions no longer wipe out per-bot overrides — with a new endpoint to do that resetting explicitly.

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

01 Migration guardrails and modes in fastly-tf-migrate NEW

95

`fastly-tf-migrate` now requires a legacy `fastly/fastly` provider version between 9.1.0 and below 20.0.0, read from `.terraform.lock.hcl`, before it analyzes or migrates `fastly_service_vcl` resources. Migrations can be previewed read-only with `fastly-tf-migrate migrate --dry-run`, which shows registry-checked HCL previews without touching state, or applied with `fastly-tf-migrate migrate --write`, which creates private recoverable HCL backups and produces stable output for CI change detection.

Verify what `fastly-tf-migrate` will change before committing — safe to run in CI or locally without touching any state.

```
$ fastly-tf-migrate migrate --dry-run
```

Apply the migration with automatic HCL backups so you can recover if something goes wrong, and get stable diff output for CI pipelines.

```
$ fastly-tf-migrate migrate --write
```

— Names exact flags, version range and lock file; runnable commands given

Fastly Terraform Provider (beta) 0.2.9 …

02 IAM role authorization for S3 backends NEW

95

Backends can connect to private AWS S3 buckets using an IAM role instead of access keys by setting `iam_vendor`, `iam_product`, `iam_role`, and `iam_region` — and optionally `iam_external_id` — on a backend create or update; Fastly automatically signs requests to the bucket.

Grant Fastly access to a private S3 bucket without storing secret keys by attaching an IAM role to a CDN service backend via the API.

```
$ curl -s -X POST
'https://api.fastly.com/service/<service_id>/version/<version_id>
/backend' \
  -H 'Fastly-Key: <api_token>' \
  -H 'Content-Type: application/x-www-form-urlencoded' \
  -d 'name=my-s3-backend' \
  -d 'address=my-bucket.s3.amazonaws.com' \
  -d 'port=443' \
  -d 'use_ssl=1' \
  -d 'iam_vendor=aws' \
  -d 'iam_product=s3' \
  -d 'iam_role=arn:aws:iam::123456789012:role/FastlyS3Access' \
  -d 'iam_region=us-east-1' \
  -d 'iam_external_id=optional-external-id'
```

— Names every config field and gives a runnable backend-creation command

Use IAM roles to connect backends to AW…

03 **IAM-authorized backend request metrics** NEW 85

Two new metrics — `origin_fetch_iam_success_count` and `origin_fetch_iam_failure_count` — are now available in the Historical Stats API and the Real-Time Analytics API to monitor IAM-authorized backend requests.

Query the Historical Stats API to confirm an IAM role is set up correctly by checking success and failure counts for IAM-authorized backend requests.

```
$ curl -s -H 'Fastly-Key: <api_token>'
'https://api.fastly.com/stats/service/<service_id>?
field=origin_fetch_iam_success_count&field=origin_fetch_iam_failu
re_count&from=2026-11-10T00:00:00Z&to=2026-11-10T23:59:59Z'
```

— Names exact metrics and APIs, includes a runnable query

Monitor IAM-authorized backend requests…

04 **Bot category actions preserve per-bot overrides** BREAKING

70

Setting a category action in the Bot Management API now only applies to bots set to `inherit`, leaving bots with their own action — including Fastly default actions on new workspaces — unchanged; it no longer resets every bot in the category. A new API endpoint sets all bots in a category to `inherit`, resetting per-bot overrides while leaving the category's own action unchanged and returning every bot in the category in the response; workflows that relied on the old bulk-reset behavior must now explicitly call this endpoint afterward.

— Explains before/after behavior and migration note, but endpoint path not given

Bot Management: category changes keep b…

---

05 SubtleCrypto deriveBits, encrypt, decrypt support NEW

60

The JavaScript SDK can now perform cryptographic operations — `deriveBits`, `encrypt`, and `decrypt` — using the `SubtleCrypto` API, now available via StarlingMonkey v0.5.0.

— Names the exact APIs and engine version but no usage example

JavaScript SDK 4.2.0 (2026-11-10)

└──▷ BREAKING ON UPGRADE

! Setting a category action no longer resets every bot in that category to `inherit` — bots with their own action now keep it. Workflows that relied on a category action change to bulk-reset all bots must now explicitly call the new bulk-inherit endpoint afterward.

WAS THIS USEFUL? 👍 👎

◆ Code Security

skylos

Feedback ⓘ v5.6.0 NOTES ↗ RANK

> Skylos shipped new `setup` and `open` CLI commands that tie together a single workflow from repo setup through to the PR comment, expanded supply-chain scanning to cover malware feeds, install hooks, agent/publish-job risks and upgrade review, and rolled out new security checks across a long list of languages plus improvements to cleanup, dependency scanning and secrets detection.

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

01 **New `setup` and `open` CLI commands** NEW 85

`skylos setup` connects AI coding agents, wiring Claude Code and Codex hooks and writing the PR check in one step; `skylos open` opens an interactive daily view showing the codebase graph, recent changes, and all check results. Together these give skylos a single workflow from setup through to the PR comment.

Bootstrap a repo for use with an AI coding agent — wires Claude Code and Codex hooks and writes the PR check in one step.

```
$ skylos setup
```

Open the interactive daily view showing the codebase graph, recent changes, and all check results.

```
$ skylos open
```

— Two named runnable commands with clear purpose v5.5.0

02 **New security checks across Ruby, Terraform, PHP, TypeScript, Swift, C#, Go, Java, Kotlin, and Python**

NEW DETECTION NEW 75

Adds Ruby and Terraform AWS declaration checks, PHP checks for command execution, code execution, XSS, and Laravel SQL injection, TypeScript NoSQL injection checks with wider request taint tracking, Swift dead-code detection alongside improved TS/JS, Go, and Rust checks, and C# modeling of ASP.NET sources, typed sinks, and unsafe deserialisation paths. Also adds CWE IDs and expanded security coverage for Python, models SQL libraries, XSS, and framework sinks in the Go engine, expands Java security

checks with Spring request tracking, and adds Kotlin security checks alongside improved Kotlin dead-code detection.

— Lists specific languages and checks but no CLI flags v5.5.0

THINNER COVERAGE BELOW

03 **Inline `skylos ignore` suppression comments** NEW 50

Suppress findings by adding `skylos ignore` comments directly in source files, supported across every language skylos checks.

— Names the comment but not its exact syntax v5.5.0

04 **Expanded dependency scanning with removed-API line references**

IMPROVED 50

Expands dependency scanning across more package managers and now names the specific lines of code that use an API a dependency upgrade removed.

— No package managers or flags named v5.5.0

05 **Scoped cleanup plans with verification and Cloud receipts**

NEW 45

skylos adds scoped cleanup plans, stronger verification, and Cloud receipts to the cleanup workflow.

— Names three additions with no commands or flags v5.5.0

06 **More credential and private key detection patterns** IMPROVED 30

skylos adds more credential and private key detection patterns in secrets scanning.

— No specific pattern names given v5.5.0

NEW DETECTION High-risk infrastructure checks

UPDATED DETECTION Static npm import evidence in SCA for JS/TS

v5.6.0 NOTES ↗ LATEST STABLE

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/duriantaco/skylos:5.6.0
```

pipx

```
$ pipx install skylos==5.6.0
# already installed? upgrade to this version:
$ pipx install --force skylos==5.6.0
```

git

```
$ git clone --branch v5.6.0 https://github.com/duriantaco/skylos.git
# already have the repo? check out this version:
$ git checkout v5.6.0
```

published under the project's GitHub account · PyPI attestation · checked 11 Nov

Wrong command?

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

GitHub App support for Enterprise Copilot seat management (2026-11-11) BLOG ↗

RANK

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

> GitHub shipped enterprise API access for Copilot seat management, added Google's Gemini 4 Carbon model to Copilot, automated MSBuild setup for the C++ language server, introduced automatic pausing of Dependabot alerts on inactive repositories, and refreshed the dashboard's default views.

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

01 Gemini 4 Carbon model in Copilot NEW

85

GitHub Copilot adds Google's Gemini 4 Carbon model for complex coding, agentic, and reasoning tasks with a focus on secure code generation. Copilot Enterprise and Business administrators can manage access via the model policy in Copilot settings (Policies > Model policy); new models are automatically enabled unless the global default is off or the model is explicitly disabled.

Switch to Gemini 4 Carbon in Copilot CLI to get security-focused code suggestions for a complex refactoring task.

> 📍 In the model picker within your chosen surface (VS Code, JetBrains, Xcode, Copilot CLI, github.com, etc.), select `Gemini 4 Carbon` from the list of available models.

Restrict or allow Gemini 4 Carbon for your organisation before it auto-enables for all users.

> 📍 In GitHub organisation settings, go to Copilot › Policies › Model policy, find Gemini 4 Carbon, and set it to `Enabled` or `Disabled` to override the global default.

— Includes exact UI navigation and model picker steps

Gemini 4 Carbon in GitHub Copilot (2026…

02 Dependabot alerts auto-pause on inactive repositories

85

BREAKING

Starting December 8, 2026, Dependabot alerts on repositories with no activity in the past year are automatically turned off and archived, with archived alerts restored on any new push. Repositories with an enforced security configuration enabling Dependabot

alerts, or in enterprises whose Dependabot alerts policy doesn't allow admins to toggle alerts, are exempt. Applies to all plans on github.com and GitHub Enterprise Cloud with data residency, where archived alerts are stored in the same region as the rest of the enterprise's data.

Re-enable Dependabot alerts on a dormant repository and restore its archived alerts without waiting for a push — navigate to the repository settings and toggle alerts back on.

> 📍 In the repository, go to Settings › Code security and analysis › Dependabot alerts and click `Enable`.

— Gives exact date, exemptions, and re-enable navigation

`Dependabot alerts will pause on inactiv…`

---

03 **Auto-configured MSBuild builds in C++ language server**

IMPROVED — 60

The MSBuild compile commands extractor is now bundled directly with the Microsoft C++ Language Server and invoked automatically on project open, so go-to-definition, find references, and other code intelligence features work without a manual or agent-driven setup skill or a hand-authored `compile_commands.json`. Compilation database generation for MSBuild projects is now deterministic, reducing setup token usage before Copilot can start working with the code.

— Explains mechanism and before/after but no user-facing command

`C++ language server now auto-configures…`

THINNER COVERAGE BELOW

04 **Enterprise App permission for Copilot seat management** NEW — 55

Enterprise GitHub Apps can now call the enterprise-level Copilot seat management APIs to assign and remove Copilot seats for users and enterprise teams, authorized via the new `Enterprise Copilot license management` fine-grained permission.

— Names the exact permission but not the API endpoints

`GitHub App support for Enterprise Copil…`

---

05 **Dashboard default view changes** IMPROVED — 40

GitHub's dashboard now shows six pull requests and issues by default instead of three, and lets you pin the Feed or Home tab as your default view. Long titles on pull request, issue, and agent session lists now stay on one line with full text available on demand, and the agent sessions section is now hidden when no agent sessions exist on the user's account.

— Describes changes concretely but no config or endpoint named

A more informative GitHub dashboard (20…

└──▷ BREAKING ON UPGRADE

! Starting December 8, 2026, Dependabot alerts on repositories with no activity in the past year are automatically turned off and archived; teams relying on passive alert monitoring for dormant repositories will stop receiving new alerts unless alerts are re-enabled.

WAS THIS USEFUL?  

OpenQodex

Feedback ⓘ openqodex@0.17.0 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 0.17.0 adds support for reviewing files and diffs of any size — oversized JSON, huge snapshot files, and 500,000+ line files — alongside new coverage reporting in reports and receipts, and three breaking changes to the library API.

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

01 **`Change.diffs[].text` now optional, use `diffHunks`** BREAKING

72

`Change.diffs[].text` is now optional — spilled entries carry a `file` path instead of inline text. Library callers that read diffs directly must switch to `diffHunks` for retained patches.

— Names exact field, replacement API and migration path openqodex@0.17.0

02 **`LineSet` required for `runScanners` coverage** BREAKING

72

Library callers must now pass `coverage` to `runScanners` as a `LineSet` built with new LineSet().addRange(start, end) — a `Map` of `Set<number>` is no longer accepted.

— Exact API call and rejected type both given openqodex@0.17.0

03 **Byte-range labels for long changed lines** NEW

65

OpenQodex now carries changed lines of any length across briefs and pages as labelled UTF-8 byte fragments, with labels showing byte ranges such as 'a.js: line 7, bytes 1 to 2,000 of 6,000'.

— Concrete label format shown but no command to trigger it openqodex@0.17.0

04 **Coverage visibility in reports and receipts** IMPROVED

62

Saved reports now name every unexamined changed range separately and show coverage by file and range in Markdown and HTML output; every review receipt prints coverage counts and names the first ten unexamined ranges, pointing to the report for the rest. The report page also shows cut rows of changed files with very long lines, such as minified bundles, when they fit once cut.

— Names output formats and receipt behaviour but no exact UI path openqodex@0.17.0

THINNER COVERAGE BELOW

05 **Unlimited-size snapshot file review** IMPROVED

51

OpenQodex now keeps snapshot files of any size in review, reading large files by indexed line ranges, streaming non-UTF-8 files byte-by-byte, and crediting only the exact bytes a reviewer reads.

— Explains mechanism but no concrete config surface openqodex@0.17.0

06 **`DIFF_CAP_BYTES` export removed** BREAKING

50

The `DIFF_CAP_BYTES` export has been removed; library callers must stop importing it.

— Clear removal with precise required action, little explanation of why openqodex@0.17.0

07 **Structural summaries for oversized JSON reviews** NEW

48

OpenQodex now reviews oversized JSON changes and dependency stores through structural summaries, keeping short changed values, record counts, outliers and store manifests visible and paginated without losing changed ranges.

— Describes behaviour but no flag or command to invoke it openqodex@0.17.0

08 **Compact interval-based coverage tracking** IMPROVED

43

Coverage now uses compact intervals instead of one entry per line, extending changed-line tracking to files with over 500,000 lines.

— Names the limit but no mechanism detail beyond that openqodex@0.17.0

└──▷ **BREAKING ON UPGRADE**

! The `DIFF_CAP_BYTES` export has been removed; library callers must stop importing it.

! `Change.diffs[].text` is now optional — spilled entries carry a `file` path instead of inline text. Library callers that read diffs directly must switch to `diffHunks` for retained patches.

! Library callers that pass `coverage` to `runScanners` must now build it with new LineSet().addRange(start, end); passing a `Map` of `Set<number>` is no longer accepted.

WAS THIS USEFUL?  

SonarSource SonarQube

Feedback ⓘ announced · Nov 10, 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 a full portfolio hierarchy management UI for Enterprise customers, letting admins build, navigate and safely restructure nested portfolio trees.

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

01 Portfolio hierarchy management UI in SonarQube Cloud NEW

82

SonarQube Cloud Enterprise plan customers get a visual portfolio hierarchy UI rendering a full tree view of parent-child portfolio relationships at unlimited nesting depth. Circular reference protection blocks invalid changes (e.g. placing a parent under its own child) before saving with a plain-language explanation; guided link and unlink flows require an explicit confirmation step before any relationship is committed; a breadcrumb trail shows the current position at every level navigated; branches can be expanded and collapsed so child portfolios load on demand for performance at scale; and the tree is fully keyboard-navigable with ARIA roles applied to parent-child relationships.

Build and navigate a multi-level portfolio hierarchy as an Enterprise admin without losing track of your position.

> 📍 In SonarQube Cloud, go to your Enterprise admin console › Portfolios › select a portfolio › open the Hierarchy view › use the breadcrumb trail to confirm your position, then choose `Link child portfolio` or `Unlink` and complete the guided confirmation flow before the change commits.

— Full UI path and mechanism given, but no API/config names. launch-20261110-722ba9c4

NEW SUPPORTED PLATFORM Intended architecture language support

NEW INTEGRATION SonarQube plugin for Devin CLI

> announced · Nov 10, 2026 BLOG ↗
>
> SonarQube Cloud adds intended architecture definition via coding agents, CLI, MCP Server, and public Web APIs for programmatic management.
>
> Install architecture skills for Claude Code or Cursor with `Sonar architecture install-skill`, or explore available commands with `Sonar architecture --help` via the SonarQube CLI.

Add `architecture` to the `SONARQUBE_TOOLSETS` configuration in the SonarQube MCP Server to enable architecture management through any MCP-compatible agent.

Define, read, replace, and delete a project's intended architecture programmatically via public Web API endpoints, with `Browse` permission for reads and `Administer architecture` permission for writes.

SonarQube Cloud now lets teams define component boundaries and allowed relationships as intended architecture, surfacing violations automatically as maintainability issues in existing pull request workflows.

`announced · Nov 10, 2026` `BLOG ↗`

SonarQube adds a Devin CLI plugin with secrets-scanning hooks, an MCP Server, and Sonar Vortex cross-file analysis inside the agent's own loop.

Install the SonarQube plugin for Devin CLI with `devin plugins install SonarSource/sonarqube-agent-plugins` to give Devin access to SonarQube skills.

Run `/sonarqube:sonar-integrate` inside a Devin session to install secrets scanning, the SonarQube MCP Server, and Sonar Vortex (when entitled) for every project at once; configuration is written to `~/.config/devin/`.

Link a repository to a SonarQube project with `sonar link <projectKey>` in the repository root so all skills resolve the correct project automatically.

List SonarQube projects from inside a Devin session with the `/sonarqube:sonar-list-projects` slash command.

Pull open issues from inside a Devin session with the `/sonarqube:sonar-list-issues` slash command.

`+4 more`

WAS THIS USEFUL? 👍 👎

syzkaller

Feedback ⓘ changes since 2026-11-10 CODE ↗

Syzkaller is an unsupervised coverage-guided kernel fuzzer developed by Google that automatically finds bugs in OS kernels by generating and executing random syscall sequences.

syzkaller adds loong64 KVM fuzzing support, including setup, register assertions, and exit tests for LoongArch64 guests

Assert KVM register state in loong64 guests using the new `syz_kvm_assert_reg` pseudo-syscall.

Run KVM exit tests against loong64 guests via new exit-test descriptions added to `sys/linux`.

Support loong64 in the `check-syzos.sh` tooling via the `loong64` branch of `$TARGETARCH`, which maps to the `loongarch64` objdump target.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ SCA updates (2026-11-10) BLOG ↗ LATEST

STABLE

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

Veracode SCA adds `--appname` and `--show-cli` agent scan flags, Unified Policy workspace assignment for EU and US Federal regions, and build-break control tied to fix availability.

Link agent scan results to an application profile by passing `--appname` to the SCA agent scan command (Commercial region; requires OAuth or HMAC authentication).

Save JSON output to a file and display a human-readable summary in the console at the same time by passing `--show-cli` together with `--json` to the SCA agent scan command; using `--show-cli` without `--json` returns an error.

Assign policies to SCA agent-based scan workspaces in the US Federal region using the Unified Policy feature to create issues and break builds based on defined criteria.

Assign policies to SCA agent-based scan workspaces in the European region using the Unified Policy feature to create issues and break builds based on defined criteria.

Configure the Vulnerability Severity policy rule (Commercial region) to break builds only when an upgrade path to fix a vulnerable open-source component is available.

+1 more

NEW SUPPORTED PLATFORM Poetry package manager support for SCA Upload scans

NEW DETECTION Malicious package detection

UPDATED COMPLIANCE FRAMEWORK Auto-Update OWASP and CWE Top 25 policy rules updated to 2025 standards

NEW COMPLIANCE FRAMEWORK Six new built-in security policies added

▾ 1 more release in this issue · 2026-11-09 → 2026-11-10

> Platform updates (2026-11-09) BLOG ↗ STABLE
>
> Veracode adds v2 Reporting API endpoints, a Custom PDF REST API, OAuth Client Credentials for XML APIs, MPT Flaw Matching, and container/IaC scan results in the platform UI.
>
> Request findings, scans, deleted scans, and audit log reports from dedicated v2 Reporting API endpoints — `POST /analytics/findings/report`, and equivalent endpoints for scans, deleted scans, and audit logs — replacing the deprecated `POST /appsec/v1/analytics/report`.

Filter findings reports by date ranges using four new optional filters on `POST /analytics/findings/report`: `first_found_start_date`, `first_found_end_date`, `resolved_start_date`, and `resolved_end_date`.

Generate custom PDF reports programmatically via the new Custom PDF REST API.

Authenticate with Veracode XML APIs using OAuth Client Credentials, in addition to the existing REST API support.

Create and manage OAuth 2.0 Client Credentials on the API Credentials page in the Veracode Platform for use with Veracode integrations and APIs.

+14 more

└─▷ **BREAKING ON UPGRADE**

- ! The `/healthcheck/status` endpoint has been removed and now returns a 404 error; use `v2/principal` in the Identity REST API instead.
- ! The `report_type` field is no longer accepted in the request body for v2 Reporting API endpoints.
- ! Success events have been removed from the Firewall dataset in Veracode Analytics.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Endor Labs

Feedback ⓘ

Introducing AI Sast That Thinks Like A Security Engineer |… | Endor Labs (2026-11-11)

BLOG ↗ RANK

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

> Endor Labs launched AI SAST in private preview — a multi-agent system that triages findings and reviews PR business logic — then extended it with first-class C support, deployment- and context-aware learning from repo files and threat models, and bring-your-own-key support for model providers.

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

01 Deployment-aware organizational context learning NEW

91

AI SAST automatically reads `AGENTS.md`, `CLAUDE.md`, skills files, and AI rules files from the repository to extract security-relevant context such as authentication middleware and user-facing scope, without configuration. It also ingests IaC files — Dockerfiles, Terraform, CloudFormation, Kubernetes YAML, and Helm charts — to build a deployment model inferring network boundaries and component interactions for every PR and local scan. Teams can upload architecture and threat model context in plain natural language, which AI SAST verifies against code and classifies into project context, true positives, and false positives, scopeable to a single project, a set of projects, or the entire tenant; reviewers can also attach thumbs-up or thumbs-down feedback to individual findings during triage to adjust severity for entire issue classes and suppress ruled-on patterns in future scans.

Suppress a class of XSS findings already mitigated by an upstream WAF by marking them as false positives during triage so future scans deprioritize that pattern automatically.

> 📍 In the AI SAST findings view, open a reflected XSS finding, click the thumbs-down icon, and add a note such as 'Mitigated by CDN-managed WAF ruleset in front of all public traffic' — repeat for the same class to train severity adjustment across future scans.

Encode organization-wide authentication standards as scan-time enforcement so AI SAST flags any new code that deviates from approved auth libraries.

> 📍 Upload a plain-language security context document scoped to the entire tenant that states your standard auth libraries and prohibited patterns; AI SAST classifies

matching deviations as true positives in every project scan.

— Named file types, categories, and exact triage steps given

How Endor Labs AI SAST Learns What Matt… | Your model, our analysis: BYOK for Endo…

02 **Bring-your-own-key for model providers** NEW 78

BYOK lets teams route all AI SAST and AI Chat model requests through their own OpenAI, Azure OpenAI, Gemini, or Vertex AI account, with provider API keys either encrypted in the Endor Labs platform or passed as environment variables at scan time so keys are never stored with Endor Labs. BYOK enforces a provider allowlist and tenant isolation so requests go only to approved providers and keys and requests remain separate from other tenants.

— Named providers and key handling, no exact config key shown

Your model, our analysis: BYOK for Endo…

03 **C language support in AI SAST** NEW 74

AI SAST now supports C with buildless scans that run in the IDE and on pull requests, covering buffer overflows, use-after-free, integer overflows, and cross-function flaws without requiring a compiler or build to reproduce. C SAST and C SCA now run together in a single pass, covering first-party C code and the open-source it depends on, and each C finding includes a call path, a working exploit, and a suggested fix — the same evidence format used for other supported languages.

— Clear scope and mechanism, no exact command shown

Why Endor Labs AI SAST for C finds what… | Your model, our analysis: BYOK for Endo…

THINNER COVERAGE BELOW

04 **AI SAST multi-agent triage and code review** NEW DETECTION NEW 56

AI SAST (private preview) automatically classifies every finding as true positive, false positive, or unknown across the full OWASP 2025 Top 10, including broken access control and insecure design, before results reach the security team. An AI Security Code Review agent assesses pull requests for architecture design and business logic flaws affecting security posture, working alongside the SAST detection and triage agents.

— Describes mechanism but no UI path or command given launch-20261111-65428615

Introducing AI Sast That Thinks Like A Security Engineer |… | Endor Labs (2026-11-11)

BLOG ↗ STABLE

Endor Labs adds AI SAST: a multi-agent system using a Code API, Opengrep, and OPA-backed policies to detect business logic flaws and auto-triage findings.

> Endor Labs now adds AI SAST — a multi-agent static analysis capability that orchestrates detection, triage, and remediation agents to surface complex vulnerabilities such as IDORs, broken access control, and insecure design flaws.
>
> Triage now runs automatically on every finding, classifying each as true positive, false positive, or unknown before it reaches the security team.
>
> A new Code API builds a semantic model of code on top of which callgraphs, dataflow tracing, and code search operations can be performed, enabling agents to reason across multi-file and multi-function interactions.
>
> AI SAST integrates with Endor Labs' policy engine, powered by Open Policy Agent (OPA), allowing teams to write custom policies to surface only high-confidence findings, suppress false positives, or block risky pull requests.
>
> Custom prompts let organizations teach the system about their own architectural patterns, security requirements, and risk priorities, with user feedback on classifications continuously refining future results.

WAS THIS USEFUL? 👍 👎

Manifest

Feedback ⓘ announced · Nov 10, 2026 BLOG ↗ RANK

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

> Manifest's single release this window adds transitive dependency mapping to the Components tab, letting users trace which direct dependency introduced any vulnerable or unexpected transitive component.

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

01 Transitive dependency mapping in Components tab NEW

79

The Components tab now labels every component as Direct, Transitive, or Unknown based on SBOM dependency graph data, showing an introduced-by path so you can see which direct dependency pulled in a given transitive component. A relationship filter lets you show Direct, Transitive, or Unknown components in any combination, with search results preserving the parent direct dependency context; this graph building works in SaaS, self-hosted, and air-gapped deployments, reading only the uploaded SBOM with no network calls, and supports both CycloneDX and SPDX formats.

When a vulnerable library makes headlines and you need to know which of your direct dependencies pulled it in, open the asset and filter the Components tab to show only Transitive components to immediately see the introduced-by path.

📍 In The Manifest Platform, open an asset, go to the Components tab, and set the relationship filter to Transitive. Expand the direct dependency listed under Introduced by to see every transitive component it brings in, with version, license, supplier, and vulnerability columns intact.

— Names exact UI location, filter states, formats, and deployment modes.

launch-20261110-c84ff51e

WAS THIS USEFUL? 👍 👎

ks-aur-scanner

Feedback ⓘ v3.0.0 NOTES ↗ LATEST STABLE

ks-aur-scanner v3.0.0 adds whole-directory scanning, obfuscated-command resolution, `--force` install override, and fail-closed stage isolation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.0.0 https://github.com/KiefStudioMA/ks-aur-scanner.git
# already have the repo? check out this version:
$ git checkout v3.0.0
```

Build with `aur-scan install --force` to bypass a confirmed finding and proceed deliberately.

NEW DETECTION Whole package directory is now scanned, not only files the PKGBUILD names — covering committed archives, binary-looking files, subdirectories, and arch-specific arrays.

NEW DETECTION Obfuscated command names (empty insertions, parameter-expansion operators, `$"..."` quoting, subcommand-prefixed options) are now resolved; unresolvable cases are flagged.

Printed text can no longer suppress a High or Critical finding.

A crash in any stage now fails closed for that package only, and previously unbounded inputs are now bounded.

+1 more

NEW DETECTION META-008 detection code for plain-literal echo lines

UPDATED DETECTION Approximately 25 new detection codes added in v3.0.0

UPDATED DETECTION PERSIST-002 timer detection

└──▷ BREAKING ON UPGRADE

! About 25 new detection codes mean packages that passed quietly under 2.2.2 may now halt — use `aur-scan install --force` after review.

! Printed-text suppression of High or Critical findings no longer works.

! The example rule no longer auto-loads; it has moved to `/usr/share/aur-scanner/example-rule.toml`.

! History keys changed; incompatible going forward (2.x records are still readable).

! `aur-scanner-core`'s `AnalysisContext` and `ParsedPkgbuild` gained new fields, breaking library users at the source level.

▾ **1 more release in this issue** · 2026-11-09 → 2026-11-10

v2.2.2 NOTES ↗ STABLE

`check` now resolves and scans AUR packages that substitute for a dependency via `provides`, noting which providers it added; `install` is unchanged.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.2.2 https://github.com/KiefStudioMA/ks-aur-scanner.git
# already have the repo? check out this version:
$ git checkout v2.2.2
```

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ

Client log exports now include metadata requests, not just downloads (2026-11-10)

BLOG ↗ RANK

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

> Cloudsmith expanded client log exports to capture the full request chain behind a download, not just the download itself.

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

01 Metadata requests added to client log exports IMPROVED

35

Client log exports now include format-specific metadata requests alongside downloads, so the full request chain behind a download is visible, rather than only the download event itself.

– Describes what changed but no export location or format details given

Client log exports now include metadata…

WAS THIS USEFUL? 👍 👎

JFrog

Feedback ⓘ

Introducing JFrog Threat Hub: Turning Vulnerability Intelligence into Blast Radius Analys

BLOG ↗ RANK

JFrog software supply chain solutions provide an end-to-end pipeline to control your binaries from build to production. Power your software updates to the edge.

> JFrog introduced Threat Hub, a new dashboard inside Xray that turns vulnerability intelligence into one-click blast radius analysis across an organization's artifacts.

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

01 **Check Impact blast radius search** NEW 90

The Check Impact button within Threat Hub runs an organization-wide impact search that lists every affected artifact with its repository, vulnerable component version, scan date, and runtime status. Results can be exported or used to set Xray policies directly from the Threat Hub interface without leaving the platform.

When a new supply-chain attack is announced, scope your blast radius immediately without cross-tool pivoting.

> 📍 In JFrog Xray, go to Xray › Overview to open Threat Hub. Find the relevant attack or CVE post in the feed, then click Check Impact to see every affected artifact, its repository, vulnerable component version, scan date, and runtime status.

Scope the blast radius of a new supply chain attack immediately after disclosure, without leaving the Xray interface.

> 📍 In JFrog Xray, go to Xray › Overview to open Threat Hub. Find the relevant CVE or malicious-package post, then click Check Impact to see every affected artifact, its repository, vulnerable component version, scan date, and runtime status across your organization.

— Exact button, output fields, and step-by-step usage given

Introducing JFrog Threat Hub: Turning V… launch-20261110-6647bd61

02 **Threat Hub dashboard in Xray** NEW 70

A new Threat Hub dashboard, accessed via `Xray > Overview`, provides a daily-updated feed of vetted CVE advisories, attack reports, and malicious package alerts covering npm, PyPI, and Hugging Face. Posts show a Severity Mismatch Indicator when

JFrog Security Research's real-world exploitability assessment differs from the CVE's published CVSS score, helping teams prioritize by actual risk.

— Names exact nav path and data sources; no limits given

Introducing JFrog Threat Hub: Turning V… launch-20261110-6647bd61

WAS THIS USEFUL? 👍 👎

zarf

Feedback ⓘ v1.0.0 NOTES ↗ RANK

Zarf is an open-source tool that enables automated, air-gapped deployment of Kubernetes applications and their dependencies into disconnected or secure environments.

> Zarf v1.0.0 adds chunked resumable OCI blob pulls, state compression, and a non-interactive `--confirm` flag for `patch-git`, alongside a major 1.0 breaking-change cleanup that makes `v1beta1` init packages the default, deprecates `v1alpha1`, and strips out legacy SDK fields, commands and server components.

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

01 **`--confirm` flag for `zarf dev patch-git`** NEW 80

Add `--confirm` to `zarf dev patch-git` to skip the interactive prompt in automation, for example `zarf dev patch-git --confirm <git-repo-path>` to run non-interactively in a CI pipeline.

Run `patch-git` non-interactively in a CI pipeline without being prompted for confirmation.

```
$ zarf dev patch-git --confirm <git-repo-path>
```

— exact flag plus a runnable example command v1.0.0

02 **v1beta1 init package becomes default, v1alpha1 deprecated** 70

BREAKING

Zarf now defaults to the `v1beta1` init package format and deprecates `v1alpha1` packages, signalling the path toward their future removal. Using variables with `v1beta1` now errors, and `v1beta1` init package values chart inputs are now limited, so previously accepted fields may be rejected.

— names exact package versions and new error/validation behavior v1.0.0

03 **Legacy CLI feature and command removal** BREAKING 65

Daemon pull-through via crane is removed, the `zarf tools registry export` subcommand is removed, the built-in SBOM viewer is removed, the artifact server is deleted, and previously deprecated commands and flags are removed.

— names exact removed subcommand and components v1.0.0

THINNER COVERAGE BELOW

04 **Go SDK breaking changes** BREAKING — 55

Deprecated SDK fields and bridge methods are removed from the Go SDK, the deprecated Kubernetes API migration helper is removed, and global architecture state is replaced with per-operation options, changing SDK and CLI call sites.

— names SDK-level changes but no migration path detail v1.0.0

05 **Hard error on package secret write failure** BREAKING — 45

Package secret write failures now cause a hard error instead of being silently ignored.

— clear before/after but no remediation steps given v1.0.0

06 **Chunked resumable OCI blob pulls** NEW — 42

OCI pulls now use chunked, retryable, resumable blob transfers to handle large packages over unreliable links.

— explains mechanism but gives no flag, size limit or command v1.0.0

07 **Zarf state field removal** BREAKING — 40

The `storageclass` and `distro` fields are removed from Zarf state.

— names removed fields but gives no migration guidance v1.0.0

08 **Zarf state compression** IMPROVED — 30

Zarf state is now compressed, reducing cluster secret size for large deployments.

— no mechanism or size figures given, thin description v1.0.0

└──▷ **BREAKING ON UPGRADE**

! The `storageclass` and `distro` fields are removed from Zarf state.

! The `v1beta1` init package is now the default; `v1alpha1` packages are deprecated and using variables with `v1beta1` now errors.

! Deprecated SDK fields and bridge methods are removed from the Go SDK.

! The deprecated Kubernetes API migration helper is removed.

! Daemon pull-through via crane is removed.

! The `zarf tools registry export` subcommand is removed.

! The built-in SBOM viewer is removed.

! The artifact server is deleted.

! Previously deprecated commands and flags are removed.

! Global architecture state is replaced with per-operation options, changing SDK and CLI call sites.

! `v1beta1` init package values chart inputs are now limited; previously accepted fields may be rejected.

! Package secret write failures now cause a hard error instead of being silently ignored.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.167.8 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 configurations.

> Infisical's v0.167.8 release adds CSR-based certificate renewal to its PKI module, a new CockroachDB account type for PAM, a bulk-expand control on the secrets overview page, and an overhauled billing UI with subscription management.

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

01 **Certificate renewal via imported CSR in PKI** NEW 45

Infisical's PKI now accepts a CSR on import so certificate renewal can reuse an existing key pair.

– Explains mechanism but no endpoint or flag named v0.167.8

02 **Open All control on secrets overview** NEW 40

The secrets overview page now has an Open All control to expand all resource types at once.

– Names exact UI control and its behavior v0.167.8

03 **Billing UI subscription management** NEW 40

Infisical's billing UI adds a subscriptions list, product details sheet, and in-sheet plan management.

– Names three UI additions without navigation detail v0.167.8

04 **CockroachDB account type in PAM** NEW 35

Infisical adds a CockroachDB account type to its privileged-access management (PAM) dynamic accounts.

– Names the integration but no config details v0.167.8

WAS THIS USEFUL? 👍 👎

◆ Application Security Posture

OX Security

Feedback ⓘ

Agentic AI Security: How OX VibeSec Proves the Value of Risk That Never Happened (2026-1

BLOG ↗

OX Security is an accelerate the transition to software that secures itself. OX Security pinpoints risk to the exact line of code, eliminating security debt from AI generation to cloud runtime.

OX VibeSec adds prompt-time secure-coding guidance, malicious dependency blocking, real-time action policies with detect/prevent toggle, and per-block OX issue records tied to application criticality.

OX VibeSec now classifies developer prompts at write time and injects the matching security guideline into the agent context, preventing vulnerable code before it is committed.

OX VibeSec now blocks malicious packages before they enter the build when a coding agent adds a dependency.

NEW INTEGRATION OX VibeSec adds real-time policy enforcement at the point of action: prompts and tool calls across CLIs, SDKs, MCPs, and skills are evaluated live and allowed, flagged, or blocked.

OX VibeSec adds a detect/prevent toggle: in detect mode the agent is alerted and the developer keeps working; flipping to prevent mode blocks the action while keeping the same OX issue as a record.

Every flag or block across all three layers opens an OX issue in the AppSec queue, showing the affected application, its criticality, the team involved, and whether the action was detected or blocked.

+1 more

NEW DETECTION Prompt injection and agent hijacking detection

NEW DETECTION Risky agent behavior detection

WAS THIS USEFUL? 👍 👎

Cycode

Feedback ⓘ

Introducing Maestro Memory: Application Security That Understands You BLOG ↗

RANK

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

> Cycode introduced Maestro Memory, giving its AI security assistant persistent context across sessions so investigations build on prior knowledge automatically.

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

01 **Maestro Memory for persistent context** NEW 55

Maestro now retains context across sessions with two new memory stores: always-on memory, which holds your preferences, rules, and team vocabulary, and recalled memory, which holds factual knowledge about products and past conversations. Maestro surfaces memory saves in-conversation and lets you search, modify, or delete what it remembers through the same conversation interface.

— Describes mechanism and two memory types but no config or API detail

Introducing Maestro Memory: Application…

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-11-09 CODE ↗

CAPEv2 is an open-source malware sandbox that automatically executes and analyzes suspicious files to extract configurations, payloads, and behavioral indicators.

CAPEv2 adds de4dotEx as the default .NET deobfuscator, replacing legacy de4dot with static .NET Reactor 6.9/7.x, ConfuserEx, and VirtualGuard support

└──▷ USE IT

Enable de4dotEx for .NET Reactor, ConfuserEx and VirtualGuard deobfuscation after upgrading, replacing the deprecated eziriz path.

custom/conf/integrations.conf

```
[de4dot_deobfuscate]
enabled = yes
binary = /usr/local/bin/de4dot

[eziriz_deobfuscate]
enabled = no
```

Enable bingraph reporting in on-demand mode now that `HAVE_BINGRAPH` is set unconditionally on package availability.

custom/conf/reporting.conf

```
[bingraph]
enabled = yes
on_demand = yes
binary = yes
procdump = yes
cape = yes
```

Switch the default .NET deobfuscator to `de4dotEx` v3.12.0 (installed at `/usr/local/bin/de4dot`) in `conf/default/integrations.conf.default`, gaining static .NET Reactor >= 6.9/7.x (CVM devirtualization, PRNG resource decryption), ConfuserEx, and VirtualGuard coverage without CoreCLR reflection.

Deprecate `[eziriz_deobfuscate]` by setting `enabled = no` in `conf/default/integrations.conf.default`, with a pointer to `[de4dot_deobfuscate]` for .NET Reactor deobfuscation.

Pin `binGraph` to the 4.0 release in `pyproject.toml` and `extra/optional_dependencies.txt`, and import it unconditionally in

`modules/reporting/bingraph.py` and `web/analysis/views.py` so `HAVE_BINGRAPH` reflects package availability regardless of `on_demand` mode.

Honor `binary`, `procdump`, and `cape` toggles under `[bingraph]` in `modules/reporting/bingraph.py`.

> └─▷ **BREAKING ON UPGRADE**
>
> ! The `[eziriz_deobfuscate]` section in `conf/default/integrations.conf.default` now defaults to `enabled = no`; existing deployments relying on `eziriz_deobfuscate` for .NET Reactor deobfuscation must migrate to `[de4dot_deobfuscate]` and install `de4dotEx`.
>
> ! The installer (`installer/cape2.sh`) now fetches `de4dotEx-3.12.0-x64.deb` from GitHub instead of the legacy `de4dot_3.1.41592.3405-2_all.deb` Ubuntu archive package; amd64 is the only supported architecture for the new package.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

reversecore_mcp

Feedback ⓘ v3.1.0 NOTES ↗ RANK

> Reversecore MCP's v3.1.0 release focuses on taming its 186-tool footprint with a token-saving router profile and a paginated evidence cache, while adding new binary-format analyzers for Go/Rust, Electron and Unity IL2CPP plus ESIL-based equivalence verification for reimplemented functions.

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

01 Router profile for reduced tool-schema token footprint NEW

95

Setting `REVERSECORE_PROFILE=router` cuts MCP tool schema token footprint by 94.4% (from ~63,000 to ~3,500 tokens) by exposing a compact 6-tool router (`search_reversecore_tools`, `call_reversecore_tool`, `get_reversecore_tool_schema`, `recommend_tools_for_goal`, `generate_analysis_plan`, `get_job_result`) while keeping all 186 tools reachable on demand. Setting `REVERSECORE_COMPACT_SCHEMA=true` additionally compresses verbose tool docstrings to single-sentence summaries under any active profile.

Cut token usage by ~94% when an agent first connects — the router profile exposes only 6 orchestration tools and loads full tool schemas only on demand.

```
$ REVERSECORE_PROFILE=router REVERSECORE_COMPACT_SCHEMA=true
docker run --rm -p 8000:8000 sjkim1127/reversecore-mcp
```

— Exact env vars, token numbers and runnable docker command given v3.1.0

02 Evidence cache for large analysis output NEW

90

get_analysis_evidence(evidence_id, offset_line, limit_lines, grep_pattern) retrieves large decompilation, disassembly, or strings output a window at a time instead of flooding LLM context, with full artifacts stored in `EvidenceCache` under an `ev_<12hex>` ID. `r2_decompile`, `r2_disassemble`, and `run_strings` now take a `detail="summary"` or `detail="full"` parameter to switch between a compact digest and full output, and `EvidenceCache` records persist across `docker run --rm` restarts via a hybrid in-memory LRU + disk-backed store written to `<workspace>/.reversecore_cache/evidence/ev_*.json`.

Decompile a large function in summary mode to get a digest without filling the LLM context, then page through the full output by grep pattern.

```
$ # Step 1 — get the summary and evidence_id
call_reversecore_tool(tool='r2_decompile', params={'address':
'0x004012a0', 'detail': 'summary'})
# -> returns evidence_id: 'ev_a3f9c1d80b2e'

# Step 2 — retrieve only the lines containing 'malloc'
get_analysis_evidence(evidence_id='ev_a3f9c1d80b2e',
offset_line=0, limit_lines=40, grep_pattern='malloc')
```

— Named function signature, flags, file path and worked example v3.1.0

03 **Electron ASAR security auditor** NEW 85

The new `unpack_and_analyze_electron_asar` tool audits an Electron app bundle for `nodeIntegration: true`, `contextIsolation: false`, `webSecurity: false`, IPC channel exposure, hardcoded API endpoints, and high-entropy secrets, requiring no external Node.js.

Audit an Electron app bundle for insecure IPC channels and hardcoded secrets during a supply-chain review.

```
$ call_reversecore_tool(tool='unpack_and_analyze_electron_asar',
params={'asar_path': '/samples/app.asar'})
```

— Concrete flags named plus a runnable call example v3.1.0

04 **Stripped Go/Rust binary recovery analyzer** NEW 80

The new `analyze_go_rust_binary` tool recovers function names, struct/interface types, build IDs, and compiler versions from stripped Go or Rust binaries, covering Go 1.2–1.20+ `pclntab` magics and both `_ZN` legacy and `_R` v0 Rust symbol formats.

— Mechanism and version coverage named, no runnable example v3.1.0

05 **ESIL-based reimplementation equivalence verifier** NEW 80

`verify_feature_equivalence` verifies behavioral equivalence of a candidate Python reimplementation against live Radare2 ESIL execution across System V AMD64 (`rdi, rsi, rdx, rcx, r8, r9`), Windows x64 (`rcx, rdx, r8, r9`), ARM64 (`x0–x5`), and ARM32 (`r0–r3`) calling conventions, reporting results in a detailed `obligation_ledger` with statuses `verified_equivalent` , `partially_verified` , or `divergent` .

— Rich mechanism and named conventions, but no runnable example v3.1.0

06 **Unity IL2CPP analyzer** NEW 75

The new `analyze_unity_il2cpp` tool recovers Unity IL2CPP string literals, metadata header offsets, and exported runtime symbols (`il2cpp_init` , `il2cpp_domain_get` , `il2cpp_runtime_invoke`), and generates ready-to-run `Il2CppDumper` / `r2` commands.

— Named symbols and outputs but no example invocation v3.1.0

THINNER COVERAGE BELOW

07 **Decompiled code to Python template reconstructor** NEW 55

The new `reconstruct_feature_code` tool reconstructs decompiled C pseudo-code into a copy-pasteable Python template with 32-bit/64-bit modular wrapping enforced via `& 0xFFFFFFFF` and `& 0xFFFFFFFFFFFFFFFF` .

— Named masks and purpose, no usage example v3.1.0

08 **Generic batch tool executor** NEW 35

The new `batch_execute_tools` generic batch executor lets repeated analysis steps run in one call.

— Named but no parameters or example given v3.1.0

WAS THIS USEFUL? 👍 👎

Finite State Platform

Feedback ⓘ

November 9, 2026 - Custom license identifiers, rescan, and Succeeded scans (2026-11-09)

BLOG ↗ RANK

Finite State Platform is a software supply chain security tool that analyzes firmware and embedded software for vulnerabilities, risks, and compliance issues.

> Finite State Platform added custom license identifier renaming, background-update tracking for False Positive status changes, expanded rescan support for older versions, and a new Succeeded count on the Scans page.

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

01 **False Positive background-update tracking** NEW 90

Checking whether findings are still updating after marking a component False Positive is now possible via `PUT /public/v0/components/{projectVersionId}/{componentId}/status`, which returns an `X-Finding-Updates-Pending` header when background updates are in progress. Components still pending such an update can be listed via `GET /public/v0/components/finding-updates/{projectVersionId}`.

After bulk-marking a large component False Positive, poll for pending finding updates before closing a review ticket.

```
$ curl -X GET
'https://platform.finitestate.io/public/v0/components/finding-
updates/<projectVersionId>' \
  -H 'Authorization: Bearer <token>'
```

— Two named endpoints with exact header and parameters

November 9, 2026 - Custom license ident…

02 **Custom license identifier renaming** NEW 88

Rename a custom license identifier, including letter-case changes, via `PATCH /v0/licenses` under Settings > Policy > Licenses. Policy rules and concluded component licenses follow the new identifier automatically after the next policy evaluation run.

Rename a misspelled or rebranded custom license identifier while keeping all policy rules and component conclusions intact.

```
$ curl -X PATCH
https://platform.finitestate.io/v0/licenses/<license-id> \
  -H 'Authorization: Bearer <token>' \
  -d '{"identifier": "MyCorrectLicense-1.0"}'
```

— Names exact endpoint, UI path, and propagation behavior

November 9, 2026 - Custom license ident…

THINNER COVERAGE BELOW

03 **Rescan support for older versions** IMPROVED 45

Rescan now works for more older versions; the first rescan of an older version prepares its original file and shows progress, while later rescans of the same version skip that preparation step.

— Describes behavior but no command or UI path given

November 9, 2026 - Custom license ident…

04 **Succeeded total on Scans page** NEW 45

The Scans page now shows a Succeeded total between the Queued and Failed counts; selecting it opens Scan History filtered to completed runs.

Filter the Scan History view to only completed scans to judge a failed-scan count in context.

> 📍 In the Finite State Platform, go to Scans, then select the Succeeded count between Queued and Failed to open Scan History filtered to runs that completed successfully.

— Names UI location and exact interaction path November 9, 2026 - Custom license ident…

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

Yamato Security Hayabusa

Feedback ⓘ v4.4.0 NOTES ↗

CODE ↗ RANK

Hayabusa is a fast Windows event log threat hunting and forensics timeline generator written in Rust, using Sigma-based detection rules.

> Hayabusa v4.4.0 adds native scanning of EWF and VHDX forensic disk images without mounting or extraction, and consolidates the pivot-keywords-list output format.

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

01 Native EWF disk image scanning NEW

93

The new `--e01` flag, available on all event-log scanning commands, lets Hayabusa scan EWF1 (`.E01`) and EWF2 (`.Ex01`) disk images natively on Linux, macOS, and Windows with no extraction, mounting, FUSE, or administrator privileges required. It handles split image segments, mixed `--vhdx --e01` directories, incomplete acquisitions, and read failures, preserving readable results and exiting nonzero on partial input.

Generate a DFIR timeline from a forensic EWF image acquired from a seized host, without extracting logs or needing admin rights.

```
$ hayabusa dfir-timeline --e01 -f disk.E01 -w -U -p verbose -o
timeline.csv
```

Scan a directory of VHDX images and E01 images together in one pass, letting Hayabusa discover all images up to three subdirectory levels deep.

```
$ hayabusa dfir-timeline --vhdx --e01 -d images/ -w -U -o
timeline.csv
```

— Flag, formats, and platform/behaviour details all given with a runnable command. v4.4.0

02 Native VHDX disk image scanning NEW

91

The new `--vhdx` flag, available on all event-log scanning commands, scans VHDX disk images natively on Linux, macOS, and Windows: a single image via `-f` or discovery of images up to three subdirectory levels deep via `-d` . Logs are read from all NTFS volumes without mounting drivers, administrator privileges, or temporary extraction; evidence paths retain the image and internal filename, and partial scans mark HTML reports incomplete and exit nonzero.

Scan a directory of VHDX images and E01 images together in one pass, letting Hayabusa discover all images up to three subdirectory levels deep.

```
$ hayabusa dfir-timeline --vhdx --e01 -d images/ -w -U -o
timeline.csv
```

— Flag, options, scope, and partial-scan behaviour all specified with example. v4.4.0

THINNER COVERAGE BELOW

03 Consolidated pivot-keywords-list output files IMPROVED

52

`pivot-keywords-list` now writes four consolidated output files — Users, IP Addresses, Processes, and Command Lines — combining fields that were previously split across three files each, and removing the logon ID and Source Computers lists.

— Names exact output files changed but no usage example given. v4.4.0

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Email Security

Sublime Security

Feedback ⓘ

Platform v1.90.208 (Agent Hub and Mean Time to Coverage) (2026-11-10) BLOG ↗

RANK

Sublime Security is an email security platform that detects and blocks phishing, malware, and other threats using customizable, logic-based detection rules.

> Sublime Security introduced Agent Hub for managing its ASA and ADÉ agents in one place, alongside a new Mean Time to Coverage metric that tracks how quickly detection gaps are closed.

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

01 **Agent Hub for ASA and ADÉ** NEW 55

A single view to manage ASA and ADÉ agents and track their impact through investigations completed, messages remediated, and estimated analyst hours saved. Accessible by going to Agents in the Sublime Security platform.

Review agent activity and your current MTTC to understand how quickly ASA and ADÉ are closing detection gaps after a missed threat is identified.

> 📍 In the Sublime Security platform, go to Agents to see investigations completed, messages remediated, estimated analyst hours saved, and your Mean Time to Coverage (MTTC).

— Names what it tracks and where to find it, not internals.

Platform v1.90.208 (Agent Hub and Mean …

02 **Mean Time to Coverage metric** NEW 50

Mean Time to Coverage (MTTC) is a new metric showing the average time from identifying a missed threat to ADÉ generating a new Detection Rule, helping teams see how quickly detection gaps are closed.

Review agent activity and your current MTTC to understand how quickly ASA and ADÉ are closing detection gaps after a missed threat is identified.

> 📍 In the Sublime Security platform, go to Agents to see investigations completed, messages remediated, estimated analyst hours saved, and your Mean Time to Coverage (MTTC).

— Explains calculation basis but no benchmarks or thresholds.

Platform v1.90.208 (Agent Hub and Mean …

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Runtime Security

Falco

Feedback ⓘ announced · Nov 11, 2026 BLOG ↗ RANK

Falco is an open-source cloud-native runtime security tool that detects anomalous behavior and threats in containers, Kubernetes, and Linux systems using kernel-level monitoring.

> Falco Operator 0.5.0 adds a central OCI artifact server, mTLS-secured artifact delivery with per-instance SPIFFE identities, a new ArtifactNode resource for rollout visibility, and compatibility enforcement that checks rules and plugins against the running Falco version before installing them.

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

01 mTLS artifact delivery with SPIFFE identities NEW

90

Falco Operator can deliver artifacts over mTLS, giving each Falco instance a unique SPIFFE identity and reloading certificates without restarts; enable it by setting `mtls.enabled=true` in the Helm chart, backed by cert-manager and trust-manager.

Enable mTLS artifact delivery in an air-gapped or security-sensitive cluster so each Falco instance proves its identity before receiving rules.

```
$ kubectl create namespace falco-operator
kubectl create namespace falco
kubectl label namespace falco-operator falco
artifact.falcosecurity.dev/trust=true
helm install falco-operator falcosecurity/falco-operator \
  --namespace falco-operator --set mtls.enabled=true
```

— Exact Helm flag, dependencies and install command given launch-20261111-77065307

02 Per-node ArtifactNode resource for rollout status NEW

90

The operator introduces the `ArtifactNode` custom resource (group `artifact.falcosecurity.dev`), one per namespace/node, whose `spec.artifacts` lists assigned resources and whose `status.artifacts` reports installed files, paths, content hashes, and conditions such as `DependenciesSatisfied`.

Inspect per-node artifact status after a rules rollout to confirm every node received and loaded the correct revision.

```
$ kubectl get artifactnodes -n falco
kubectl get artifactnode worker-1 -n falco -o yaml
```

— Named CRD fields with exact kubectl inspection commands launch-20261111-77065307

03 **Rule and plugin compatibility enforcement** NEW 90

Artifact specs can declare `required_engine_version` and `required_plugin_versions` inside Rulesfile specs, controlled by `enforceRequirements: true` (default) or `enforceRequirements: false` for advise-only mode; the operator checks compatibility against the running Falco version via the `/versions` API before installing any rulesfile or plugin, keeping Falco on the last compatible set when a new revision fails checks.

Pin a custom rulesfile to minimum engine and plugin versions so the operator withholds it from any node that cannot satisfy the requirements.

yaml

```yaml
apiVersion: artifact.falcosecurity.dev/v1alpha1
kind: Rulesfile
metadata:
  name: custom-rules
spec:
  inlineRules:
    - required_engine_version: 0.62.0
    - required_plugin_versions:
        - name: container
          version: 0.7.0
    - rule: Terminal shell in container
```

— Config keys, API endpoint and fallback behavior all named launch-20261111-77065307

04 **Central artifact server for OCI rule/plugin delivery** NEW 65

A central artifact server now runs inside the cluster on port `8082`; the operator pulls each OCI artifact once, pins it to a digest, and serves it to every Falco pod instead of each pod pulling independently from the registry.

— Mechanism and port named but no direct usage command launch-20261111-77065307

THINNER COVERAGE BELOW

05 **Client-side API rate limits for large clusters** NEW 45

Client-side API rate limits are now on by default: 50 QPS (burst 100) for the operator and 5 QPS (burst 10) for each sidecar, preparing the operator for large clusters.

— Concrete numbers given but no config key to adjust them launch-20261111-77065307

06 **Breaking changes in 0.5.0 upgrade path** BREAKING 25

This release contains breaking changes; consult the v0.5.0 migration guide before upgrading from any prior release, including 0.5.0 release candidates.

— Points to migration guide without naming specific breaks launch-20261111-77065307

└──▷ **BREAKING ON UPGRADE**

! This release contains breaking changes — consult the v0.5.0 migration guide before upgrading from any prior release, including 0.5.0 release candidates.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ changelog · Nov 11, 2026 CHANGELOG ↗ RANK

Plerion is a cloud security platform that provides comprehensive risk management, vulnerability detection, and threat prioritization across cloud environments and workloads.

> Plerion's latest release adds an Access tab that visualizes AWS resource exposure across policy layers, Bank Negara Malaysia RMiT compliance mapping, bulk continuous-scanning enablement for AWS accounts, a new AKS end-of-support check, and expanded inventory coverage for AWS service families.

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

01 **Bank Negara Malaysia RMiT compliance mapping** NEW 75

Plerion adds Bank Negara Malaysia RMiT compliance mapping, covering 96 of 472 paragraphs across AWS, Azure, GCP, and Kubernetes in areas including cryptography, network resilience, access control, cloud services, cybersecurity operations, and data security. Enable it from Compliance under Manage frameworks.

Turn on Bank Negara Malaysia RMiT as a compliance framework to map your cloud posture against the BNM policy document across AWS, Azure, GCP, and Kubernetes.

> 📍 In the Plerion console: go to Manage › Compliance › Manage frameworks › locate Bank Negara Malaysia RMiT › toggle it on.

— Specific coverage numbers and a clear enablement path changelog-20261111-a91e6e2d

02 **Access tab for AWS assets** NEW 70

A new Access tab on AWS assets lets you inspect who can reach a sensitive resource (e.g. an S3 bucket or KMS key) without writing policy queries. Graph or List views show the identity policy, permission boundary, SCP, RCP, and resource policy layers on each relationship path when you hover over an arrow.

Review who can access a sensitive S3 bucket or KMS key without writing policy queries — open the asset and switch to the Access tab.

> 📍 In the Plerion console: open any AWS asset (e.g. an S3 bucket or KMS key) › click the Access tab › choose Graph or List view › hover over a relationship arrow to inspect the identity policy, permission boundary, SCP, RCP, and resource policy layers on that path.

— Clear UI navigation but no API or config surface named changelog-20261111-a91e6e2d

THINNER COVERAGE BELOW

03 **Bulk continuous scanning enablement for AWS accounts** NEW

55

Continuous scanning can now be enabled for multiple AWS accounts at once from the account list by selecting accounts and saving; the editor shows which accounts already have it on and which still need a CDR integration.

Enable continuous scanning across a fleet of AWS accounts in one operation rather than account by account.

> 📍 In the Plerion console: go to the AWS account list › click Continuous scanning › tick the target accounts › Save — the editor flags which accounts already have it on and which still need a CDR integration.

— Exact UI steps given, no API or flag named changelog-20261111-a91e6e2d

04 **AKS end-of-support check** NEW

40

Plerion adds an AKS cluster Kubernetes end-of-support check (`PLERION-AZURE-464`) for Azure Kubernetes Service.

— Named check ID but no usage or trigger detail changelog-20261111-a91e6e2d

05 **Inventory coverage for new AWS service families** NEW

20

Plerion expanded inventory coverage to twelve new AWS service families.

— No names of the services added, mentioned only in summary changelog-20261111-a91e6e2d

WAS THIS USEFUL? 👍 👎

◆ SaaS Security Posture

Nudge Security

Feedback ⓘ

Filter browser extensions by AI agent verdict NOTES ↗

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

Nudge Security now exposes AI agent verdicts in the browser extension table, with filtering to cross-reference verdict against approval status.

UPDATED DETECTION Browser Extension Risk Analyst — AI agent verdicts surfaced in extension table

WAS THIS USEFUL? 👍 👎

Obsidian Security

Feedback ⓘ announced · Nov 10, 2026

PRESS ↗

Obsidian Security delivers complete SaaS security—gain control, stop threats, and ensure compliance across all your business apps.

Obsidian adds a Runtime Policy Decision Engine, MCP Security Proxy, and Policy Decision API to intercept and govern AI agent actions at the moment they execute.

NEW INTEGRATION New Runtime Policy Decision Engine evaluates agent actions against centrally managed policies using identity, HR, and application context, then allows, blocks, or routes for human approval.

NEW INTEGRATION New MCP Security Proxy intercepts every agent tool call before it reaches a business application, keeping app tokens out of AI clients and enforcing the same central policies used elsewhere.

New Policy Decision API lets customer-managed gateways, agent harnesses, and workflows query Obsidian for an allow/block/approve decision with reason, without rebuilding policies per gateway.

Customer-provided context — such as data classification labels — can now be supplied so runtime decisions reflect the business's own sensitivity definitions alongside identity and application data.

NEW INTEGRATION Discovery and governance for in-app agents in Salesforce Agentforce, Databricks, GitHub, Microsoft 365, Snowflake, ServiceNow, and Workday

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

calico

Feedback ⓘ announced · Nov 9, 2026 BLOG ↗

Calico is an open-source networking and network security solution for Kubernetes and other container orchestration platforms, providing policy-based traffic control.

Calico adds Red Hat OpenShift 5 support via the Tigera Operator, with eBPF/iptables/nftables data plane choice, KubeVirt live migration IP stability, and layer 2 VLAN bridging for VM workloads in Calico Enterprise.

NEW SUPPORTED PLATFORM Calico now supports Red Hat OpenShift 5, installable through the Tigera Operator listed in the Red Hat Ecosystem Catalog, with a documented migration path from OVN-Kubernetes to Calico for existing clusters.

Calico runs on eBPF, iptables, or nftables on Red Hat OpenShift 5, letting teams choose the data plane per cluster.

Calico IPAM now preserves VM IP addresses through KubeVirt live migration on Red Hat OpenShift Virtualization and advertises them via BGP into the data center network.

Calico Enterprise adds layer 2 bridging on Red Hat OpenShift 5 to extend existing VLANs into the cluster, allowing VMs to migrate to Red Hat OpenShift Virtualization without a network redesign.

WAS THIS USEFUL? 👍 👎

Chainguard

Feedback ⓘ Week of 2026-11-09 NOTES ↗ RANK

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

> Chainguard shipped SCIM 2.0 identity provisioning, a breaking-change `cgb` package endpoint, private Go and Java package upload support, and a Guarded Entrypoint startup wrapper for Custom Assembly images.

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

01 cgb breaking-change package endpoint

BREAKING

81

A new `cgb` package endpoint is introduced for packages that require an API-incompatible update, such as a CVE fix, replacing `cgp`. Repository or pinning configs can be switched from `https://cgp.chainguard.dev/...` to `https://cgb.chainguard.dev/...`, after which OSV and VEX feeds recognise `.cgb` versions as fixed.

Switch a Chainguard Athena image to the breaking-change endpoint when a CVE fix requires an API-incompatible package update.

```
📍 # In your package pinning or repository config, replace the cgp endpoint with cgb: # Before: https://cgp.chainguard.dev/ ... # After: https://cgb.chainguard.dev/ ... # OSV and VEX feeds will then recognise .cgb versions as fixed.
```

— Exact endpoint names and migration steps given. Week of 2026-11-09

02 Guarded Entrypoint secrets and preflight wrapper NEW

78

Guarded Entrypoint now adds a Chainguard-built startup wrapper to Custom Assembly images that can resolve environment variable secrets from Google Secret Manager, Vault, or Consul; build a Java truststore at runtime; run TCP or file preflight checks; and override the application command — all without requiring a derived image.

— Clear mechanism and integrations, but no exact command shown. Week of 2026-11-09

THINNER COVERAGE BELOW

03 **SCIM 2.0 user provisioning** NEW 55

Chainguard now supports SCIM 2.0, allowing admins to provision and deactivate Chainguard users directly from Okta or Microsoft Entra ID, and to map identity provider groups to built-in or custom Chainguard roles.

— Names providers and mapping but no setup steps or config key. Week of 2026-11-09

04 **Private Go and Java package uploads** NEW 23

Chainguard added support for uploading private Go and Java packages.

— Only mentioned in summary with no mechanism or endpoint. Week of 2026-11-09

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ Now Available: Trials read API NOTES ↗ RANK

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

> Huntress shipped a new Trials read API for accounts and resellers, plus delivery status tracking and resend for bounced portal invitations.

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

01 Invitation delivery status and resend on Pending Users NEW

65

The Pending Users page now shows delivery status (`Sending` , `Delivered` , `Bounced` , or `Blocked`) for account and organization invitations, with hover tooltips explaining each status and the recommended next step. Bounced invitations can be resent directly from the portal after confirming the recipient's mailbox is ready.

When an invited user never received their invite, check their delivery status and resend if bounced.

> 📍 1. Go to the Huntress portal and navigate to Pending Users. 2. Locate the invitation and check its delivery status column (Sending, Delivered, Bounced, or Blocked). 3. Hover over the status badge to read its meaning and the recommended next step. 4. If the status is `Bounced` , confirm the recipient's mailbox is ready, then click Resend.

– Names exact statuses and gives clear UI navigation steps

Resend bounced invitations from the por…

THINNER COVERAGE BELOW

02 Trials read API for accounts and resellers NEW

45

A new Trials API adds endpoints for accounts and resellers to list active and expired product trials, enabling API users to view trial statuses and select which product trials to start when creating accounts through the API.

– Describes capability but no endpoint paths or methods given

Now Available: Trials read API

WAS THIS USEFUL? 👍 👎

osctrl

Feedback ⓘ v0.6.3 NOTES ↗

osctrl is a fast, efficient open-source TLS endpoint and management platform for osquery, enabling fleet management and log collection at scale.

osctrl adds granular RBAC roles and sensitive node/carve-path controls, a configurable URL prefix, an optional DB read replica, and a `--prepare` provisioning mode.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.6.3 https://github.com/jmpsec/osctrl.git
# already have the repo? check out this version:
$ git checkout v0.6.3
```

Add a `--prepare` mode to `provision.sh` to stage environments before full provisioning.

Configure a URL prefix for the frontend and API so osctrl can be served under a subpath.

Configure an optional DB read replica so dashboard and node list reads bypass the primary database.

Route staleness-tolerant reads in `osctrl-tls` to the DB read replica to reduce primary load.

Add global RBAC roles: user management, environment management, read-only, and alert admin.

+10 more

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ Pages are generally available (2026-11-11)

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 brought Pages to general availability with new permissions and a REST API, expanded Unity Gateway with Claude Fable 5.2, and shipped several Lakeflow Designer improvements for AI Function output handling and Python preview detection.

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

01 Preview detection in Lakeflow Designer Python operators NEW

85

Lakeflow Designer Python operators can now detect whether they are running as a preview or a full run by reading `config['is_preview']`, allowing code to skip side effects like writes or API calls during previews, e.g. if not config['is_preview']: write_to_external_system(df) (Public Preview).

Guard irreversible side effects (writes, API calls) so they are skipped when a designer is previewing the operator rather than running the full pipeline.

```python
if not config['is_preview']:
    write_to_external_system(df)
```

— Exact config key and runnable code example given

Distinguish Lakeflow Designer Python pr…

02 Cleaned AI Function outputs in Lakeflow Designer NEW

65

A new Clean step for supported AI Function outputs in Lakeflow Designer splits results into separate columns for answer values, error messages, and raw responses, enabling downstream transformations (Public Preview). In the designer, select an AI Function transformation node and choose `Clean` to split the output before passing it downstream.

When building a pipeline in Lakeflow Designer, use the Clean step on an AI Function output to split results into separate columns before passing them to downstream transformations.

📍 In Lakeflow Designer, select an AI Function transformation node, then choose `Clean` to split the output into answer, error message, and raw response columns for downstream use.

— Names exact UI step and output columns, preview status

Use cleaned AI answers in Lakeflow Desi…

03 **Pages general availability with governance controls** NEW

60

Pages — governed business-concept definitions organized by domain on the Discover page — are now generally available. Admins can control who can publish and manage Pages in a domain with the new `MANAGE PAGES` permission, and Pages can now be managed programmatically via a REST API.

— Names permission and REST API but no endpoint paths

Pages are generally available (2026-11-…

THINNER COVERAGE BELOW

04 **Default external access to Unity Catalog managed Delta tables**

IMPROVED

55

For eligible workspaces with the compliance security profile enabled, external Delta clients can now create, read, and write Unity Catalog managed Delta tables by default, controlled by metastore external data access settings and Unity Catalog privileges.

— Explains scope and controls but no config key names

External access to Unity Catalog manage…

05 **Claude Fable 5.2 on Unity Gateway** NEW

35

Anthropic Claude Fable 5.2 is now available on Unity Gateway and accessible through Foundation Model APIs.

— Names model and API surface but no usage detail

Anthropic Claude Fable 5.2 is now avail…

WAS THIS USEFUL? 👍 👎

ClickHouse RunReveal

Feedback ⓘ announced · Nov 9, 2026

BLOG ↗

Transform security logs into actionable insights with RunReveal's modern SIEM platform. Unify detection, investigation, data lake, and AI capabilities at predictable costs.

RunReveal adds ClickHouse Security, an agentic SIEM in private preview deploying onto ClickHouse Cloud with AI agents, MCP server, CLI, API, and Slack bot.

Access ClickHouse Security via MCP server, API, CLI, or Slack bot to build custom agentic security workflows with the LLM of your choice.

NEW DETECTION Manage custom detections as code using the CLI — supporting `export`, `test`, `sync`, and `dry run` operations on scheduled SQL or streaming Sigma rules.

AI chat lets security team members ask follow-up investigation questions in natural language without writing traditional queries.

Built-in AI agents support alert triage, detection tuning and generation, and threat hunting without requiring a separate add-on.

NEW INTEGRATION ClickHouse Cloud integration

└──▷ BREAKING ON UPGRADE

! RunReveal's multi-tenant ClickHouse cluster option will be removed; existing RunReveal SaaS customers will be migrated to their own independent ClickHouse Cloud service.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Flashpoint

Feedback ⓘ announced · Nov 10, 2026 BLOG ↗ RANK

Flashpoint empowers the world's leading organizations with the best data for threat intelligence to protect people, places, and assets.

> Flashpoint launched Agentic AI, introducing Agentic Hub with a conversational Chat agent and five purpose-built agents that automate briefing generation, alert configuration, alert triage, PIR drafting, and Echosec report scheduling.

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

01 Agentic Hub Chat for cross-domain investigation NEW

70

Flashpoint adds a Chat agent inside Agentic Hub that turns a single natural-language question into a multi-step investigation across vulnerability, credential, fraud, and threat actor intelligence, returning source-linked findings.

When a new vulnerability is disclosed, use Agentic Hub Chat to pursue leads across exploit discussions, threat actor activity, and credential exposure in one conversational session.

📍 In Flashpoint Ignite, open Agentic Hub › Chat and type a natural-language question such as 'What threat actors are discussing CVE-2026-XXXX and is exploit code circulating?' — the agent follows related leads across vulnerability, fraud, and threat actor intelligence and returns source-linked findings.

— Clear mechanism and exact UI path with sample query given `launch-20261110-6f1f0801`

02 Alert Triage Agent for alert pre-filtering NEW

65

Flashpoint adds an Alert Triage Agent that pre-filters low-context alert results, flags alerts not matching monitoring intent, and labels and archives remaining alerts.

Reduce analyst time spent on noisy alerts by letting the Alert Triage Agent pre-filter, flag, and archive low-context results before the team reviews them.

📍 In Flashpoint Ignite, open Agentic Hub, select Alert Triage Agent, and assign it to an existing alert rule — the agent flags alerts that do not match monitoring intent and labels and archives the remainder, logging each action taken.

— Exact UI path and action sequence given in example `launch-20261110-6f1f0801`

03 PIR Agent for requirement drafting NEW

65

Flashpoint adds a PIR Agent that translates plain-language intelligence requirements into structured Priority Intelligence Requirements for team review and approval, and surfaces emerging threats not addressed by existing PIRs.

Convert a team's plain-language monitoring brief into a structured PIR ready for analyst review, without manual formatting.

> 📍 In Flashpoint Ignite, open Agentic Hub › PIR Agent and enter a plain-language requirement such as 'Track ransomware groups targeting US healthcare in Q4 2026' — the agent produces a structured PIR for team review and approval and highlights gaps in existing PIRs.

— UI path and sample requirement input provided `launch-20261110-6f1f0801`

---

04 **Intelligence Briefing Agent for role-tailored briefings** NEW

60

Flashpoint adds an Intelligence Briefing Agent that generates scheduled intelligence briefings tailored to executives, analysts, and technical teams, drawing on selected alert rules, PIRs, investigations, assets, and credential exposure data.

— Named data sources but no usage walkthrough given `launch-20261110-6f1f0801`

05 **Echosec Reporting Agent for saved-search reports** NEW

60

Flashpoint adds an Echosec Reporting Agent that transforms saved searches into scheduled reports containing assessments, cited findings, and recommended actions, preserving each report to track how a monitored situation develops over time.

— Mechanism and purpose clear but no usage steps shown `launch-20261110-6f1f0801`

THINNER COVERAGE BELOW

06 **Alert Setup Agent for alert rule creation** NEW 45

Flashpoint adds an Alert Setup Agent that converts plain-language monitoring requirements into configured alert rules, with a preview step before saving.

— Mechanism described but no concrete usage example `launch-20261110-6f1f0801`

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

oxo

Feedback ⓘ v2.22.11 NOTES ↗ RANK

> oxo extended its agent communication protocol with several new protobuf message types and fields spanning mobile, domain, IP and reputation fingerprinting, and added full certificate-chain validation to X.509 fingerprint messages.

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

01 New message types and fields across oxo's scan protocol NEW

60

oxo added several new protobuf schema elements to its agent messaging protocol: AppCards build fields in mobile file metadata messages (surfacing additional build-time metadata for mobile asset scans), store privacy data fields in the `v3.fingerprint.store` definitions (capturing app store privacy metadata during fingerprinting), the `v3.fingerprint.domain_name.service.sitemap` message type (for exchanging sitemap fingerprint data on domain name targets), the `v3.report.reputation` message type (added to the agent messaging protocol), the `v3.fingerprint.ip` network message type, and an `asn_cidr` field added to IP whois messages (exposing the CIDR block associated with an ASN).

– Names every protobuf message and field but gives no usage mechanism v2.22.10 v2.22.9 v2.22.8 v2.22.7 v2.22.6

THINNER COVERAGE BELOW

02 Certificate chain validation for X.509 fingerprint messages

IMPROVED

45

oxo now performs chain validation on X.509 fingerprint messages, extending certificate checks beyond the leaf certificate to the full certificate chain.

– Describes mechanism but no command or config to invoke it v2.22.11

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ

What's New in Iru: Korean, Remote Shell, and Auto App Digests NOTES ↗ RANK

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

> Iru's latest update adds Korean language support across the product, Remote Shell access for EDR customers, and a consolidated daily digest for Auto Apps update notifications.

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

01 **Remote Shell for EDR customers** NEW 45

Iru adds Remote Shell access for all customers using Iru EDR.



— Names the feature and eligible customer base, no further detail

What's New in Iru: Korean, Remote Shell…

02 **Daily digest for Auto Apps updates** IMPROVED 40

Iru now consolidates Auto Apps update notifications into a single daily digest.

— Describes the change but not configuration or delivery mechanism

What's New in Iru: Korean, Remote Shell…

03 **Korean language support** NEW

25

Iru now supports Korean across every product.

— States the addition but no scope or mechanism detail

What's New in Iru: Korean, Remote Shell…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Cyera

Feedback ⓘ announced · Nov 10, 2026 BLOG ↗ RANK

Cyera helps organizations protect data, govern identities, and secure AI agents to reduce risk, stay in control, and adopt AI with confidence.

> Cyera's biggest addition this window is Cy Agents, a coordinated suite of AI agents (Cy Detect, Cy Triage, Cy Learn, Cy Investigate, Cy Assist, and upcoming Cy Remediate) built into the platform, alongside Cyera Endpoint, an on-device DLP sensor for laptops and workstations, and an Omni DLP update that aggregates third-party alerts.

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

01 **Cyera Endpoint on-device DLP sensor** NEW 81

Cyera Endpoint now monitors and blocks data movement across device channels — clipboard, USB and removable media, printing, network drives, and sync clients — using data classification and business context from the Cyera platform. A local small language model runs on the device to classify content by meaning (catching sensitive data matching no pattern), Cy Triage performs real-time validation of endpoint actions and enforces blocks through the endpoint sensor before data leaves the device, policies draw on data classes already identified by Cyera DSPM, and Cy Learn captures team feedback as memories to reduce false positives over time. It covers supported native AI applications and local agents, including unsanctioned coding assistants, keeps enforcing locally for offline laptops, and is available as an add-on for customers already using Cyera Omni DLP, Cyera Agent Guardian, or both.

— Names channels, mechanism and add-on terms but no config keys launch-20261110-52e4dccb launch-20261110-de7eb3cc

02 **Cy Agents AI Agent Suite** NEW 65

(new) Cy Agents is a coordinated team of specialized AI agents directed through Cy's natural-language interface to surface exposure, prioritize risk, investigate incidents, and produce reports, with built-in audit trails keeping human teams in command of agent actions. Within the suite: Cy Detect reasons across data sensitivity, access, and exposure in combination to surface a holistic risk view instead of a flat list of findings, identifying the highest-priority exposures and assembling a full risk story around each affected asset; Cy Triage analyzes alerts from connected tools — including endpoint sources — pulling together data, identity, destination, and business context to evaluate true versus false positives, and automatically sets severity, priority, ownership, and next steps for each alert or finding; Cy Learn captures an organization's corrections, definitions, and risk tolerance as lasting platform knowledge, accepting user-defined data classes, refinements to existing classifications, and business-specific terminology,

back-applying that context to past findings and classifications, and also captures team feedback as memories that inform endpoint policies to reduce false positives over time; Cy Investigate runs end-to-end investigations — connecting identity, activity, and data context — reconstructing incidents by correlating data movement, access, and identity behavior, and delivers a findings report with evidence and stated limits on its conclusions; Cy Assist accepts natural-language queries across data, identities, alerts, and AI activity, turning plain-language requests into dashboards and reports, and supports saving and sharing prompts, scheduling recurring reports, and routing deeper work to specialized agent capabilities through supported workflows; and Cy Remediate, described as coming soon, will plan and track remediations, account for business disruption, and carry approved actions through to resolution.

— Names the agent roster and interface but no concrete controls Explains reasoning approach but gives no interface or command Clear mechanism but no exact UI or API path given Describes mechanism well but no configurable surface named Mechanism is clear; no concrete report format or path named Names concrete actions (save, schedule, route) a reader can try Announced but not yet available; no mechanism detail

launch-20261110-52e4dccb launch-20261110-ffb633a9 launch-20261110-de7eb3cc

THINNER COVERAGE BELOW

03 **Omni DLP third-party alert aggregation** NEW 35

Cyera adds to Omni DLP the ability to aggregate native endpoint alerts together with signals from third-party DLP tools, with Cy Triage prioritizing what needs attention.

— Single-line description with no mechanism detail launch-20261110-52e4dccb

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

OpenVPN

Feedback ⓘ announced · Nov 10, 2026 BLOG ↗ RANK

OpenVPN is an open-source VPN software that creates encrypted, secure tunnels for private network communications over the internet.

> OpenVPN Access Server now offers hourly billing on Azure Marketplace, giving a third purchase option alongside existing monthly/annual and BYOL plans.

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

01 **Hourly billing for Access Server on Azure** NEW 65

Access Server adds an hourly-billed listing on Azure Marketplace alongside the existing monthly/annual and BYOL options. The VM is billed per hour only while running, so stopping the VM pauses the software charge, making it suited for dev/test environments, contractor access, project migrations, and disaster-recovery standby scenarios.

Deploy Access Server on Azure with hourly billing for a short-term contractor engagement, paying only for the hours the VM runs.

📍 1. Go to Access Server on Azure Marketplace and select Get it now. 2. Select your Subscription, Resource group, and the connection tier that fits your use case. 3. Select Review + create, then Create. 4. Follow the onboarding wizard to configure client internet or private Azure network access. 5. Install OpenVPN Connect on client devices and connect with the included profile. 6. When the engagement ends, stop the VM to pause the hourly software charge, or delete the resource group to remove all Azure resources.

— Names billing mechanism and use cases with deploy steps, no pricing figures

launch-20261110-2d814b20

WAS THIS USEFUL? 👍 👎

amneziawg-installer

Feedback ⓘ v6.3.0 NOTES ↗ RANK

AmneziaWG-installer is a script that installs AmneziaWG, a modified WireGuard VPN with obfuscation features to bypass deep packet inspection and censorship.

> amneziawg-installer's two releases add live network controls — a separate exit-IP path and a port-change command that regenerates all client configs — plus PresharedKey-by-default, randomized install ports, Raspberry Pi OS trixie module packages, and a fix for Oracle Cloud Ubuntu's leftover iptables block.

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

01 **Live port-change command for running servers** NEW 90

`manage port <number>` rewrites `ListenPort` in `awg0.conf` and the init file, updates UFW rules, and regenerates every client config and QR code in one step, with automatic rollback if the server fails to come up on the new port.

Move a live server to a new UDP port and have every client config and QR code regenerated atomically, with automatic rollback if the server fails to come up.

```
$ sudo bash ./manage port 31337
```

— Exact command, config files touched, and rollback behavior v6.3.0

02 **Prebuilt kernel modules for Raspberry Pi OS trixie** NEW 90

Installs pre-built AmneziaWG module packages on Raspberry Pi OS trixie (Debian 13, kernel 6.18.60+rpt) for Pi 3/4 (64-bit), Pi 5, and the 32-bit `rpi-v7` kernel, so a `--protocol=2.0` install no longer falls back to building via DKMS on the device; the installer selects the target automatically by OS version (13 → trixie, else bookworm).

Deploy an AmneziaWG 2.0 server on a Raspberry Pi running Raspberry Pi OS trixie — the installer now selects a pre-built trixie module package instead of compiling via DKMS.

```
$ sudo bash ./install_amneziawg.sh --protocol=2.0
```

— Exact kernel versions, device targets, and install command given v6.2.2

03 **Exit-IP routing for client traffic** NEW 85

Route client traffic out through a second public IPv4 using `--exit-ip=<IPv4>` at install time or `manage exit-ip <IPv4|off>` on a running server, without changing client configs.

Switch client egress to a second public IP mid-operation — no client reconfiguration needed — to escape a blocked exit address.

```
$ sudo bash ./manage exit-ip 203.0.113.42
```

Install with a fixed exit IP and a specific entry port so clients connect on one address and their traffic leaves on another from the start.

```
$ sudo bash ./install_amneziawg.sh --exit-ip=203.0.113.42 --port=31337
```

— Exact flags and commands given, mechanism explained v6.3.0

04 **PresharedKey enabled by default for new clients** NEW

73

Every new client now gets a `PresharedKey` by default; opt out per client with `add --no-psk`, or disable it server-wide at install time with `--no-psk` (saved as `NO_PSK=1` in the init file), and override per client with `add --psk`.

— Flags and config key named but no example v6.3.0

05 **Oracle Cloud Ubuntu iptables fix** IMPROVED

73

Clears the blocking `REJECT --reject-with icmp-host-prohibited` iptables rule left in memory by Oracle Cloud Ubuntu images after UFW setup, so the VPN works immediately after install without an extra reboot; image-saved rules in `/etc/iptables/rules.v4` are cleaned up, while administrator-added rules are left untouched with a warning and the removal command.

— Exact rule and file path named, no example command shown v6.2.2

06 **Randomized default port for new installs** IMPROVED

63

New installs without `--port` now pick a random free UDP port in the range 10000–59999, excluding the WireGuard defaults `51820` and `51821`, instead of the fixed `39743/udp` that fingerprinted every installation.

— Exact port range and excluded values given, no example shown v6.3.0

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ announced · Nov 10, 2026 BLOG ↗ RANK

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

> Tailscale's Aperture introduced team-wide skill sharing, adding multiple ways to import skills and a full Skills over MCP implementation for distributing and managing skills across connected agents.

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

01 Skill import via file upload, GitHub sync, and well-known discovery

NEW 75

Skills can be imported into Aperture by uploading a `SKILL.md` file, a `.zip` archive, or pasting Markdown directly. Skills can also be synced from a public GitHub repository by pasting the repo link into Aperture, or imported from third-party providers that support the `.well-known` skill discovery specification by supplying a URL.

Import a third-party skill set from a provider that publishes a `.well-known` skill discovery endpoint, so your team gets curated skills without manual copy-paste.

> 📍 In Aperture, go to Skills › Add, choose `Use a URL`, and paste the provider domain (e.g. stripe.com). Aperture queries the `.well-known` skill discovery specification endpoint and presents a list of available skills to import.

Sync an existing bundle of skills from a public GitHub repository into Aperture so the whole team stays current without manual updates.

> 📍 In Aperture, go to Skills › Add, choose `Sync a GitHub repo`, and paste the public repository URL. Aperture imports all skills found in the repo.

— Names exact import formats and discovery spec, with UI steps `launch-20261110-e1f6627e`

02 Aperture as MCP server with skill management tools NEW

70

Aperture can be added as an MCP server to an agent, and users can ask the agent to add skills to Aperture through MCP. New MCP tools added to the Aperture MCP gateway, including a pre-built Skill Creator skill, let any agent connected to Aperture's MCP endpoint create, update, and delete skills.

Use an agent connected to Aperture's MCP endpoint to create a new skill without leaving your agent interface.

📍 In your agent (connected to Aperture's MCP endpoint), prompt: 'Using Aperture, create a new skill for summarising pull request diffs.' The Skill Creator skill's MCP tools handle creation and store it in Aperture.

— Names MCP endpoint and Skill Creator tool, with usage example `launch-20261110-e1f6627e`

THINNER COVERAGE BELOW

03 **Cross-agent LLM session visibility via MCP** NEW 45

Aperture now exposes MCP endpoints that list and read all LLM sessions across all connected agents, enabling cross-agent skill improvement.

— Describes capability but no concrete endpoint name or steps `launch-20261110-e1f6627e`

04 **Skills over MCP spec for identity-based distribution** NEW 45

Aperture now implements the Skills over MCP spec, letting admins distribute skills to individuals or teams based on identity, and letting users distribute skills across their own agents or computers.

— Names the spec but lacks a runnable or UI-level detail `launch-20261110-e1f6627e`

NEW INTEGRATION Grok Bot OAuth app integration

announced · Nov 9, 2026 BLOG ↗

Tailscale OAuth apps are now generally available, letting agents join a tailnet under the authorizing user's identity via standard OAuth 2.0, with posture-based access controls.

Tailscale OAuth apps are now generally available, letting third-party agents and integrations join a tailnet via a standard OAuth 2.0 flow as a device tied to the authorizing user's identity — so access policies, audit logs, and device quotas apply as they would for any user-owned device.

Use `custom:grokbot == true` as a device posture attribute and `srcPosture` in grant rules to give AI agent devices access to specific resources (such as `tag:mcp-servers`) while blocking broader network access inherited from the user's identity.

Use `defaultSrcPosture` in tailnet policy to ensure existing access rules do not automatically extend to agent devices, requiring an explicit `srcPosture` override for any grant you want agents to match.

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback ⓘ

11/10 November 10 – Web Application and API Update (2026-11-10) BLOG ↗ RANK

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

NEW DETECTION AppAware Library Expansion

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ Windows GUI 1.5.21 (2026-11-10) BLOG ↗ RANK

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

> Firezone's Windows GUI client gained managed-config auto-start on sign-in and a small settings UX improvement showing when a value has saved.

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

01 **Managed-config auto-start on Windows sign-in** NEW 47

The Windows client can now auto-start when a user signs in, enabled by deploying a managed config to the device.

– Describes mechanism but no exact config key or deployment path

Windows GUI 1.5.21 (2026-11-10)

02 **Visual save indicator in settings** IMPROVED 18

Settings fields now show a visual indicator when a value is saved.

– Bare UI change with no detail on scope or mechanism Windows GUI 1.5.21 (2026-11-10)

WAS THIS USEFUL? 👍 👎

◆ Backup & Recovery

Commvault Cloud

Feedback ⓘ changelog · Nov 10, 2026 BLOG ↗

Commvault Cloud is a backup, recovery, and cyber-resilience platform for hybrid, multi-cloud, and SaaS workloads.

Commvault adds AgentRecover for AI agent full-stack recovery, Recovery Studio for runbook-driven business restoration, and Data Activate for governed data flow into AI platforms.

NEW INTEGRATION Commvault adds AgentRecover, a capability to discover AI agents, map the data, applications, tools, identities, configurations, and infrastructure they interact with, audit their activity, and recover affected resources or the broader AI system to a trusted point — with Microsoft Foundry targeted for early access in November and Amazon Bedrock AgentCore targeted for December.

Commvault adds Recovery Studio, a runbook-driven workflow engine that orchestrates recovery of applications, identities, configurations, infrastructure dependencies, and datasets together to restore business operations rather than individual data items.

NEW INTEGRATION Commvault adds Data Activate for Databricks, a joint solution using built-in policy governance and automation to bring enterprise backup data into Databricks for business insights, AI applications, and agents, with Commvault also serving as a certified Databricks partner providing data protection for the Databricks Data + AI Platform.

NEW INTEGRATION Strategic partnership with Zscaler for agentic activity recovery

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Ping Identity

Feedback ⓘ 09 Nov 2026 (2026-11-10) BLOG ↗

RANK

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

> Ping Identity's latest rapid-channel release adds named push profiles for multi-brand MFA, per-realm/client control over token-exchange `may_act` enforcement, and performance-focused v2.0 journey nodes, alongside smaller UI and SDK refinements.

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

01 **Named push profiles for multi-brand MFA** NEW 90

Admins can create named push profiles in the Push Notification service to support multiple brands or mobile applications within a single realm, then select them in journeys using the new v2.0 Push Sender, Push Registration, Combined MFA Registration, and Recovery Code Collector Decision nodes; selecting the Legacy Configuration profile migrates from previous realm-level settings.

When running multiple branded mobile apps in one realm, create a dedicated push profile for each brand so journeys can target the correct push sender.

📍 In the admin console, go to Native Consoles › Access Management › Realm › Services › Push Notification Service, create a named push profile, then open your authentication journey and configure the Push Sender v2.0 node to select that profile.

— Names all new nodes and migration path, with concrete admin steps

09 Nov 2026 (2026-11-10)

02 **Per-realm may_act enforcement control for token exchange** NEW 80

Admins can disable `may_act` enforcement for token exchange per realm or per client, enabling subject token exchange from third-party identity providers or authorization servers while still applying equivalent authorization checks.

Allow federation partners using a third-party authorization server to exchange subject tokens without `may_act` enforcement, while keeping your own authorization checks in place.

📍 In the admin console, go to Native Consoles › Access Management › Realm › Services › OAuth2 Provider (or the specific client), and disable the `may_act`

enforcement toggle for token exchange.

— Named field and clear console path to enable it 09 Nov 2026 (2026-11-10)

THINNER COVERAGE BELOW

03 **Performance-optimized v2.0 journey nodes** IMPROVED 50

Version 2.0 of the Identify Existing User node reduces data requested during user lookup, and version 2.0 of the Query Filter Decision node reduces data requested during filter evaluation, both improving journey performance.

— Names nodes but no detail on mechanism or magnitude 09 Nov 2026 (2026-11-10)

04 **Checkmark indicators for password requirements** IMPROVED 30

The default registration journey now uses checkmarks and Xs to indicate which password requirements are met, replacing the previous color-only indicators.

— Clear UI change but no config or navigation detail 09 Nov 2026 (2026-11-10)

05 **Updated PingOne Signals (Protect) SDK** IMPROVED 20

The PingOne Signals (Protect) SDK is updated to provide the latest client-side data collection for PingOne Protect.

— Only states an SDK update with no specifics 09 Nov 2026 (2026-11-10)

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Network, Ory Hydra, Ory Kratos, Ory Keto v26.2.13 released NOTES ↗ RANK

Discover the latest features, improvements, and updates to Ory.

> Ory shipped three Enterprise-focused additions this window: custom OAuth token prefixes, per-provider SAML audience storage, and configurable session risk rules.

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

01 **Per-provider SAML sp_entity_id storage** NEW 60

Ory Network can now store the Service Provider entity ID (the audience sent by the IdP) per SAML provider using the `sp_entity_id` field.

— Names exact field but lacks usage example Ory Network, Ory Hydra, Ory Kratos, Ory…

THINNER COVERAGE BELOW

02 **Custom OAuth 2.0 token prefix per project** NEW 40

Enterprise projects on Ory Network can now set a custom prefix for issued OAuth 2.0 access tokens on a per-project basis.

— Names the capability but no config key or mechanism given

Ory Network, Ory Hydra, Ory Kratos, Ory…

03 **Session risk rule configuration** NEW 35

Enterprise plan projects on Ory Network can configure session risk rules via Ory Kratos Identities evaluation.

— Describes capability but no config keys or mechanism detail

Ory Network, Ory Kratos, Ory Polis v26.…

WAS THIS USEFUL? 👍 👎

RIG Security Feedback ⓘ

Secretless Isn't Riskless: Introducing SPIFFE Posture NOTES ↗

RIG Security is a protect human, non human, and AI identities in one platform. Map attack paths, enforce least privilege, and stop identity threats.

RIG Security adds SPIFFE Posture: agentless, read-only audit of workload-identity trust policies with cross-platform correlation and Entra ID remediation plans.

NEW DETECTION RIG Security now correlates matching SPIFFE issuer and subject pairs across all supported integrations to surface the combined reach of a single workload identity across multiple platforms.

RIG Security now generates executable remediation plans for supported Entra ID findings, letting teams review and apply changes — such as removing federated credentials or stripping directory roles — directly from the finding, with a revert path and verification against Entra ID.

NEW DETECTION SPIFFE Posture checks — AWS

NEW DETECTION SPIFFE Posture checks — privileged federation

NEW INTEGRATION SPIFFE Posture platform integrations

WAS THIS USEFUL? 👍 👎

◆ Privileged Access Management

Teleport

Feedback ⓘ announced · Nov 10, 2026 BLOG ↗ RANK

Teleport is an open-source access platform that provides secure, identity-based access to SSH servers, Kubernetes clusters, databases, and web applications.

> Teleport adds Phases to its Beams runtime, bringing Git-like snapshot, fork, rebase, and rebuild management to container filesystems, plus a new way to run persistent background agent tasks inside a beam.

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

01 Phases: Git-like filesystem management for Beams NEW

98

Teleport's Beams runtime adds Phases, a container filesystem feature with snapshot, fork, rebase, and periodic rebuild capabilities, defined in a Kube-style `phase.yaml` file supporting `version`, `metadata.name`, `from`, `files`, `commands`, and `services` fields. Create a phase interactively with `tsh phase add <name>`, fork an existing beam's filesystem state with `tsh phase add <name> --fork`, apply a definition with `tsh phase apply -f phase.yaml`, and save or finalize the current `/home` state with `tsh phase snapshot` or `tsh phase snapshot --final`. The `services` block starts a background or pre-start service inside a phase using `type: pre-start` or `type: background`, scoped to `context: user` or `context: snapshot`.



Fork a running beam to give a parallel agent its own isolated filesystem and have it commit work independently before saving a snapshot.

```
$ tsh phase add beak-buttons --fork
beamctl start sh -c 'codex exec -C ~/parakeets "Add a beak-
friendly interface. Test it, commit, push, then run tsh phase
snapshot." </dev/null'
```

Define a reproducible dev environment with a local file overlay, an apt+Go install rebuild stage, and a background dev server that starts automatically on each beam launch.

```yaml
version: v1
metadata:
  name: dev-files
files:
  - from: ./dev-files
    to: /opt/dev
---
version: v1
metadata:
  name: environment
from: dev-files
commands:
  - shell: |
      apt-get update
      apt-get install -y curl
  - shell: |
      v=$(curl -fsSL 'https://go.dev/dl/?mode=json' | jq -r '.
[].version | select(startswith("go1.27."))')
      curl -fsSL "https://go.dev/dl/$v.linux-amd64.tar.gz" | tar
-xz -C /usr/local
      ln -sf /usr/local/go/bin/go /usr/local/bin/go
services:
  - name: dev-server
    type: background
    context: snapshot
    shell: |
      export GOCACHE=/tmp/go-build GOPATH=/tmp/go
      go run /home/beams/parakeets/main.go
```

Apply a phase definition and open an interactive beam on top of the resulting environment.

```
$ tsh phase apply -f phase.yaml
tsh phase add dev --from=environment
```

— Full command set, flags, and config schema given with runnable examples.

launch-20261110-ca529eff

THINNER COVERAGE BELOW

02 **Persistent agent tasks via beamctl start** NEW 50

Run a persistent agent task inside a beam without staying connected using `beamctl start <command>`.

Fork a running beam to give a parallel agent its own isolated filesystem and have it commit work independently before saving a snapshot.

```
$ tsh phase add beak-buttons --fork
beamctl start sh -c 'codex exec -C ~/parakeets "Add a beak-friendly interface. Test it, commit, push, then run tsh phase snapshot." </dev/null'
```

— Single command named but no further mechanism or scope detail. launch-20261110-ca529eff

NEW DETECTION MITRE ATT&CK session scoring for AI agent behavior

> announced · Nov 10, 2026 BLOG ↗
>
> Teleport adds Beams, a secure runtime for AI agents with micro-VM isolation, delegated X.509 identities, phases for state snapshotting, MCP tool-level RBAC, and JIT agentic access requests.
>
> Run AI agents in isolated micro-VMs that auto-expire in 12 hours using `tsh beams add --role=<scoped-role>` to assign a least-privilege delegated X.509 identity scoped to the task instead of inheriting a user's full role set.
>
> Snapshot, fork, and branch a running agent's filesystem state using `tsh phases add`, `tsh phase snapshot`, and `tsh phases add --fork` — giving agents consistent state across runs without rebuilding from scratch.
>
> Create agent-initiated just-in-time access requests for elevated, time-boxed access using `tsh request create --resource=<resource-id> --reason=<reason>`, reviewable by a human or another agent and auto-expiring on task completion.
>
> Connect to MCP servers with tool-level RBAC and short-lived certificates instead of shared API keys using `tsh mcp login <mcp-server-name>`, scoped to only the tools the developer is allowed to use.
>
> Capture every message sent to an inference model — including LLM traces, tool-use, and protocol-level traffic — in a full audit trail for every agentic session.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 9th, 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 added VirusTotal and expanded Sumo Logic and ServiceNow integrations, introduced a live Active Syncs panel with sync cancellation, and split the Reports experience into a read-only Viewer and editable Designer.

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

01 **VirusTotal integration with watch lists** NEW 80

Add a VirusTotal integration from the Integrations page using an API key and a watch list of domains, public IP addresses, URLs, and file hashes, with entries validated as you type.

Connect VirusTotal to monitor suspicious domains and file hashes as part of your compliance evidence collection.

📍 In Compyl, go to Integrations › VirusTotal, enter your API key, then add domains, public IP addresses, URLs, or file hashes to the watch list — entries are validated as you type.

— Names exact UI path, auth method, and watch-list types Nov 9th, 2026

02 **Active Syncs panel and sync cancellation** NEW 75

A new Active Syncs panel on the Integrations page shows running, queued, stale, and cancelling syncs with live timers. Syncs can be cancelled mid-flight, stopping at the next saved step while already-pulled rows are preserved.

Stop a long-running sync mid-flight without losing already-pulled rows — useful when a large asset inventory sync is consuming resources.

📍 In Compyl, go to Integrations › Active Syncs, locate the running sync, and click Cancel — the sync stops at the next saved step and preserved rows remain.

— Gives exact cancel steps and sync-state names Nov 9th, 2026

03 **Reports Viewer and Designer split** IMPROVED 65

Reports now open in a read-only Viewer by default, with a header toggle to switch to the Designer; admins get separate links for each, and old report designer bookmarks redirect automatically.

— Explains default behavior, toggle, and redirect mechanism Nov 9th, 2026

THINNER COVERAGE BELOW

04 **Sumo Logic audit event evidence** IMPROVED 55

The Sumo Logic integration now pulls audit events so controls can show evidence of admin activity like role changes, access key creation, and 2-step verification changes.

— Names event types but no setup steps given Nov 9th, 2026

05 **ServiceNow password policy and SSO sync** IMPROVED 55

The ServiceNow integration now syncs password policy rules and SAML/SSO identity provider settings, with each endpoint validated independently.

— Names synced settings but lacks usage steps Nov 9th, 2026

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 4.2.0 (2026-11-11) BLOG ↗ RANK

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

> Runlayer's 4.2.0 and 4.1.0 releases span MCP-based agent configuration and version history, expanded AI Watch audit logging, a PyPI-installable CLI, private GitHub skill imports, an LLM Gateway devcontainer feature, all-users/all-agents policies, Terraform provider changes for agent ownership and secret handling, Microsoft Teams bot integration, and several smaller collaboration and infrastructure fixes.

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

01 RunlayerInstance Redis connections now enforce TLS

BREAKING

95

A `RunlayerInstance` with `spec.redis.tls: true` and `credentials.passwordFrom` previously connected without TLS; it now connects over TLS. If your Redis does not serve TLS, set `tls: false` explicitly.

— Exact before/after behavior and migration step given 4.1.0 (2026-11-11)

02 LLM Gateway support in devcontainers

NEW

90

Use LLM Gateway inside devcontainers by adding the `ghcr.io/runlayer/devcontainer-features/llm-gateway` Feature, or by running `runlayer credential setup` in your image and mounting `~/.runlayer/aiwatch`; usage rules, budgets, and offboarding apply as on the host.

Run Claude Code or Codex inside a devcontainer and route all LLM calls through the developer's own gateway — with usage rules and budgets enforced just as on the host.

```
$ runlayer credential setup
```

Add the LLM Gateway devcontainer feature to a devcontainer so containerised coding agents inherit the host AI Watch gateway configuration automatically.

```json
{
  "features": {
    "ghcr.io/runlayer/devcontainer-features/llm-gateway": {}
  }
}
```

— Exact feature name, command, and mount path given `4.1.0 (2026-11-11)`

03 **Terraform agent ownership transfer and sensitive server outputs**

90

BREAKING

Setting `owner_id` on `runlayer_agent` transfers agent ownership, shown in `terraform plan` with the same checks as an in-app ownership transfer, and the `runlayer_server` resource and data source now hide `transport_config` headers, `env`, placeholder defaults, and `test_placeholder_values`, requiring `sensitive = true` on root module outputs referencing these values or the plan errors.

Transfer ownership of a Runlayer agent via Terraform so governance stays in code and shows a reviewable diff in `terraform plan`.

```hcl
resource "runlayer_agent" "my_agent" {
  name     = "my-agent"
  owner_id = "usr_abc123"
}
```

— Exact attributes and breaking migration requirement given `4.1.0 (2026-11-11)`

04 **Private GitHub skill imports** NEW

75

Import skills from private GitHub repos by pasting a repo URL in Add skill → GitHub; Runlayer reads the repo through the workspace's GitHub App connection with a short-lived read-only token, then treats the result as a regular workspace-wide skill with auto-update.

— Clear UI path and token mechanism described `4.1.0 (2026-11-11)`

05 **Guest edit permissions on calendar events** NEW

75

Set guest permissions on calendar events with `create_event` and `update_event`; pass `edit_as_guest: true` to `update_event` on events you do not organize to save changes on the organizer's calendar when guest edits are allowed.

— Exact field and usage condition given `4.1.0 (2026-11-11)`

06 **Runlayer CLI installable from PyPI** NEW 70

Install the Runlayer CLI from PyPI with `uv tool install runlayer` on Python 3.14.

Install or upgrade the Runlayer CLI on a Python 3.14 environment to run agent and policy commands.

```
$ uv tool install runlayer
```

— Exact runnable install command given 4.2.0 (2026-11-11)

07 **Super Admin user identity rebind via API** NEW 65

Super Admins can move an existing user to a new email address and/or directory identity via the API, preserving roles, teams, and history, with every rebind audit-logged.

— Names API capability but not the exact endpoint 4.2.0 (2026-11-11)

08 **AI Watch Enforce blocks unlisted MCP servers** NEW 60

AI Watch Enforce now makes and records block decisions for MCP servers not found in a client's configuration, surfacing calls in Shadow AI Activity and the audit log as an unidentified MCP server instead of blocking silently on the device.

— Clear before/after behavior, no reader-facing action 4.2.0 (2026-11-11)

09 **Audit logging expanded for every checked MCP call** IMPROVED 60

Every MCP call that AI Watch Protect or Enforce checks now gets an audit log entry, including blocked unreadable requests, error-state calls, calls from deactivated-user devices, and anonymous device calls recorded against the organization.

— Names four logged call types, logging is automatic 4.2.0 (2026-11-11)

10 **All-users and all-agents policy targeting** BREAKING 60

Target every user or every agent account — including ones added later — with a single policy using the new `all users` and `all agents` principals; self-hosted deployments must run 4.1.0 or later on every backend instance before saving one, or older instances will skip these policies silently.

— Principal names and version prerequisite given 4.1.0 (2026-11-11)

11 **Scoped access for offloaded large tool outputs** IMPROVED 60

Restrict offloaded large tool outputs fetched with `get_large_output` to only the user and agent grant that produced them, even when multiple users delegate to the same agent.

— Names exact function, behavior is automatic 4.1.0 (2026-11-11)

12 **Microsoft Teams bot integration** NEW 60

Add agent bots and the Runlayer Teams notification bot to Microsoft Teams shared and private channels via Manage channel → Apps; already-installed bots update automatically.

— Clear UI path, limited mechanism detail 4.1.0 (2026-11-11)

THINNER COVERAGE BELOW

13 **Agent management via MCP** NEW 55

Agents can now be configured and managed directly from your MCP client: configure network access, environment variables, and collaboration settings, and list, read, and restore agent versions (restores are saved as new versions so they can be undone).

— Describes scope but no exact MCP command shown Describes mechanism, no exact MCP call shown

4.2.0 (2026-11-11) 4.1.0 (2026-11-11)

14 **Faster agent startup on OpenCode harness** IMPROVED 50

Agents start approximately 9 seconds sooner on a fresh sandbox on the OpenCode harness, and no longer stall for over a minute when their network policy blocks the npm registry.

— Concrete metrics given, no user action needed 4.2.0 (2026-11-11)

15 **Copy user access link for agent accounts** NEW 50

Copy a user access link from the agent account page with the new Copy user access link button so regular users can grant an agent account access on their behalf.

— Exact button name and location given 4.1.0 (2026-11-11)

16 **Managed tunnel connections auto-scale capacity** IMPROVED 50

Managed tunnel connections now automatically add capacity as stream usage grows and remove idle endpoints after their streams finish, with each endpoint defaulting to 256 concurrent streams.

— Concrete default given, scaling is automatic 4.1.0 (2026-11-11)

17 **Terraform import guide for existing objects** NEW 40

A new guide explains how to bring servers, policies, and other app-created objects under Terraform using import blocks and generated configuration.

— Points to a guide, no exact commands given `4.2.0 (2026-11-11)`

└──▷ BREAKING ON UPGRADE

! Root module Terraform outputs that reference `transport_config` headers, `env`, placeholder defaults, or `test_placeholder_values` on `runlayer_server` now require `sensitive = true` or the plan will error.

! Self-hosted deployments must run 4.1.0 or later on every backend instance before saving an `all users` or `all agents` policy; older instances skip these policies silently.

! A `RunlayerInstance` with `spec.redis.tls: true` and `credentials.passwordFrom` previously connected without TLS; it now connects over TLS. If your Redis does not serve TLS, set `tls: false` explicitly.

WAS THIS USEFUL? 👍 👎

◆ Agent Sandboxes

NVIDIA OpenShell

Feedback ⓘ v0.1.8 NOTES ↗ RANK

OpenShell runs autonomous AI agents in policy-controlled sandboxes with isolated filesystem, network, and credential access.

> OpenShell's v0.1.8 release pairs new observability and Kubernetes scheduling controls with a substantial set of breaking changes spanning credential placement, WebSocket handling, OCSF event schemas, sandbox filesystem policy, and SDK distribution.

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

01 Distributed tracing propagation and OCSF correlation NEW

95

The CLI now forwards `TRACEPARENT` and `TRACESTATE` from its environment to the gateway so CI pipeline traces parent gateway spans; OCSF HTTP and API audit events now correlate with sampled OpenTelemetry traces, and OTLP log records carry trace and span context.

Carry a CI trace context into gateway spans so the full pipeline appears as one distributed trace.

```
$ export TRACEPARENT=00-0af7651916cd43dd8448eb211c80319c-
b7ad6b7169203331-01
openshell sandbox create --name ci-run -- python run_agent.py
```

— Named env vars plus a runnable CLI example v0.1.8

02 Explicit credential placement policy for provider auth

90

BREAKING

Static provider credentials without placement metadata now resolve only in the `Authorization` header or through explicitly enabled body/WebSocket rewrite policies; custom headers such as `x-api-key`, query parameters, and path slots now require an explicit placement declaration. `auth_style` selects the active placement, and fields for other styles are ignored; operators must upgrade the gateway before the supervisor and update existing imported profiles.

— Full before/after mechanism with named fields and migration steps v0.1.8

03 WebSocket protocol enforcement on REST endpoints BREAKING

80

REST endpoints now return `403` for HTTP upgrade requests in both enforce and audit modes. WebSocket traffic must be configured with `protocol: websocket` and explicit handshake and message rules, and `websocket_credential_rewrite: true` now requires a WebSocket endpoint.

— Named config keys and exact response code, no full example v0.1.8

04 **OCSF sandbox attribution field changes** BREAKING

80

Sandbox attribution in OCSF output has moved: supervisor events now use `device.container.uid` and `device.container.name`; gateway API Activity and Detection Finding events use `resources` entries with `type: "sandbox"`; other gateway classes use `unmapped.sandbox_id` and `unmapped.sandbox_name`. SIEM parsers reading the old top-level `container.uid` or `container.name` fields must update their mappings.

— Exact old and new field names across three event classes v0.1.8

05 **Kubernetes and Helm pod scheduling controls** NEW

75

Set `supervisor_runtime_class_name` in the Kubernetes driver config to assign a runtime class to dedicated supervisor pods, and set `priorityClassName` and `topologySpreadConstraints` in the Helm chart to control gateway pod scheduling.

Pin dedicated supervisor pods to a hardware-isolated runtime class in Kubernetes for stronger workload separation.

toml

```toml
[openshell.drivers.kubernetes]
supervisor_runtime_class_name = "kata-containers"
```

— Exact config keys with a working TOML example v0.1.8

06 **IPv4-mapped IPv6 address evaluation for TCP policy** BREAKING

70

Transparent TCP connections now evaluate IPv4-mapped IPv6 destinations as IPv4, so endpoint rules such as `host: "::ffff:203.0.113.7"` must be changed to `host: "203.0.113.7"`.

— Exact before/after config values given v0.1.8

07 **TLS and mTLS for the metrics endpoint** NEW

65

The `/metrics` listener can now be served over TLS or mutual TLS, configured through gateway TOML, CLI flags, environment variables, or Helm Secrets; plaintext remains the default if nothing is configured.

– Names config surfaces but not exact flag or env var names v0.1.8

08 **Go SDK shares CLI OIDC token file** BREAKING 65

The Go SDK now shares `oidc_token.json` with the CLI; run `openshell gateway login` once to replace token files written by earlier Go SDK releases.

– Exact file name and migration command given v0.1.8

09 **Go SDK module path deprecated** DEPRECATED 60

The Go SDK module `github.com/NVIDIA/OpenShell/sdk/go` is deprecated in favor of `github.com/ai-openshell/sdk-go`.

– Exact old and new module paths, no further migration detail v0.1.8

10 **Python SDK distribution moved to PyPI only** BREAKING 60

The `openshell` Python SDK wheel and checksum are no longer attached to core GitHub Releases; install it via `pip install openshell` from PyPI.

– Runnable install command, but no reasoning given v0.1.8

THINNER COVERAGE BELOW

11 **Symlink enforcement in sandbox filesystem policy** BREAKING 55

Sandboxes now fail to launch when a filesystem policy path goes through a symlink inside a writable workspace, including in best-effort mode; operators should use the symlink's target path in the policy instead.

– Clear fix instruction but no exact policy syntax shown v0.1.8

12 **Sticky session routing for shared-supervisor pools** IMPROVED 40

In multi-replica gateways, shared-supervisor pool sandbox sessions now stay pinned to the gateway replica that serves their supervisor, including across failures.

– Explains mechanism but no named config to act on v0.1.8

13 **Supervisor route advertisement for gateway discovery** NEW 25

Supervisors now advertise their installed application routes so gateways can discover them, for both dedicated and shared supervisors.

– Describes behaviour only, no config or command given v0.1.8

└──▷ BREAKING ON UPGRADE

! Static provider credentials without placement metadata now resolve only in the `Authorization` header or through explicitly enabled body/WebSocket rewrite policies; custom headers such as `x-api-key`, query parameters, and path slots require an explicit placement declaration. `auth_style` selects the active placement and fields for other styles are ignored. Upgrade the gateway before the supervisor and update existing imported profiles.

! REST endpoints now return `403` for HTTP upgrade requests in both enforce and audit modes. Configure WebSocket traffic with `protocol: websocket` and explicit handshake and message rules; `websocket_credential_rewrite: true` now requires a WebSocket endpoint.

! OCSF sandbox attribution moved: supervisor events use `device.container.uid` and `device.container.name`; gateway API Activity and Detection Finding events use `resources` entries with `type: "sandbox"`; other gateway classes use `unmapped.sandbox_id` and `unmapped.sandbox_name`. SIEM parsers reading top-level `container.uid` or `container.name` must update their mappings.

! Sandboxes now fail to launch when a filesystem policy path goes through a symlink inside a writable workspace, including in best-effort mode. Use the symlink's target path in the policy instead.

! Transparent TCP connections now evaluate IPv4-mapped IPv6 destinations as IPv4. Change endpoint rules such as `host: "::ffff:203.0.113.7"` to `host: "203.0.113.7"`.

! The Go SDK module `github.com/NVIDIA/OpenShell/sdk/go` is deprecated; switch to `github.com/ai-openshell/sdk-go`.

! The `openshell` Python SDK wheel and checksum are no longer attached to core GitHub Releases; install via `pip install openshell` from PyPI.

! The Go SDK now shares `oidc_token.json` with the CLI. Run `openshell gateway login` once to replace token files written by earlier Go SDK releases.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

ToolHive

Feedback ⓘ v0.64.0 NOTES ↗ RANK

> ToolHive exposes its embedded auth server's TLS listener through Kubernetes Services, with a caveat that affected workloads restart once to pick up the change.

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

01 Embedded auth server TLS listener exposed via Service

78

BREAKING

When `embeddedAuthServer.tlsListener` is enabled, MCPServer, MCPRemoteProxy, and VirtualMCPServer workloads now get an `https-auth` container port and a Service port 8443 (`appProtocol: https`), making the auth server reachable through the Kubernetes Service — and, for `serviceType: LoadBalancer` or `NodePort` VirtualMCPServers, reachable from outside the cluster. As a result, affected pods are rolled once on the next reconcile to pick up the new `https-auth` container port and Service port 8443.

— Names exact config key, port, protocol, and rollout impact v0.64.0

└──▷ **BREAKING ON UPGRADE**

! MCPServer, MCPRemoteProxy, and VirtualMCPServer pods are rolled once on the next reconcile when `embeddedAuthServer.tlsListener` is enabled, due to the addition of the `https-auth` container port and Service port 8443.

WAS THIS USEFUL? 👍 👎

tenuo

Feedback ⓘ v0.3.6 NOTES ↗ RANK

> Tenuo v0.3.6 adds opt-in single-use (consume-on-use) approvals and per-call warrant resolvers for Google ADK, tightens argument- and chain-based authorization across FastAPI, Temporal, A2A and several agent frameworks, and ships a cluster of breaking changes affecting custom nonce backends, the LangGraph extra, policy enforcement semantics, and the `@tenuo/mcp` npm package.

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

01 **Single-use approval consumption** NEW 87

Opt-in consume-on-use approvals can be enabled by calling tenuo.enable_approval_consumption(), or by providing an enforcement-point store with atomic reservations and owner checks; approved calls expose an idempotency key for downstream deduplication.

Enable single-use approvals in a shared deployment so each approved call can only be redeemed once, with the idempotency key available for downstream deduplication.

```python
import tenuo

tenuo.enable_approval_consumption()
```

— Names the exact function, mechanism and a runnable example v0.3.6

02 **Stricter requirements for custom nonce backends** BREAKING 78

Custom Python nonce backends must now expose immutable retention, PoP bucket-size, and clock-skew settings, atomic admission, and owner-aware reservation and release; new stores default to 210 seconds retention including a 60-second clock budget, and namespaces must be configured before serving.

— Lists exact required properties and default values for migration v0.3.6

03 **LangGraph extra version and middleware requirements** BREAKING 75

The LangGraph extra now requires `langgraph>=1.0` and `langgraph-prebuilt>=1.0.2` on Python 3.10+, and uses `TenuoMiddleware` for `create_agent` graphs.

— Exact version pins and middleware class readers must adopt v0.3.6

04 **Per-call warrant resolver for Google ADK** NEW 68

Google ADK now supports trusted per-call warrant resolvers via `warrant_resolver`; missing warrants or resolver failures deny calls even when local allow rules are present, and resolver errors do not expose backend secrets.

— Names the resolver hook and deny behaviour but no code sample v0.3.6

THINNER COVERAGE BELOW

05 **A2A delegation chain validation and replay admission** IMPROVED 58

A2A now validates the complete delegation chain before replay admission and honors configured nonce stores, coordinating approval and replay admission while preserving denial evidence through receipts, audit logs, and observe mode.

— Describes mechanism and evidence surfaces but no config keys v0.3.6

06 **Authorization hardening across framework integrations**

IMPROVED 55

FastAPI and Temporal now authorize the arguments handlers actually receive rather than the declared schema, with JSON types preserved across adapters; CrewAI now rejects adapter agents that bypass tool hooks; OpenAI handoffs preserve restrictive policies and require constrained arguments for guarded calls; and LangChain and LangGraph retain authorization at execution boundaries.

— Explains before/after behaviour without a config or flag Names four frameworks' behaviour but no APIs or flags

v0.3.6

07 **Policy enforcement semantics breaking changes** BREAKING 50

Numeric and negated-policy enforcement points must be upgraded together, as mixed versions may disagree on value comparison semantics; conditional CEL error guards must now be kept inside the CEL expression when conditional evaluation is required, since externally guarded patterns no longer work the same way.

— States two migration risks without exact config or flag names v0.3.6

08 **npm package version pin update** BREAKING 45

`@tenuo/mcp` (npm) now requires `@tenuo/core >=0.3.6-beta.0 <0.3.7`.

— Exact version constraint but no explanation of impact v0.3.6

09 Parent chain in delegated Python warrants IMPROVED

30

Delegated Python warrants now carry the parent chain needed by root-trusting verifiers.

— Single bare sentence with no mechanism or usage detail v0.3.6

└──▷ BREAKING ON UPGRADE

! Upgrade numeric and negated-policy enforcement points together — mixed versions may disagree on value comparison semantics.

! Keep conditional CEL error guards inside the CEL expression when conditional evaluation is required; externally guarded patterns no longer work the same way.

! Custom Python nonce backends must now expose immutable retention, PoP bucket-size, and clock-skew settings, atomic admission, and owner-aware reservation and release, and namespaces must be configured before serving.

! The LangGraph extra now requires `langgraph>=1.0` and `langgraph-prebuilt>=1.0.2` on Python 3.10+.

! `@tenuo/mcp` (npm) now requires `@tenuo/core >=0.3.6-beta.0 <0.3.7`.

WAS THIS USEFUL?  

◆ AI Runtime Defense & Guardrails

pipelock

Feedback ⓘ v3.8.0 NOTES ↗ RANK

> pipelock v3.8.0 adds nearby-keyword gating for custom DLP patterns, broadens encoded-text decoding for DLP and injection scanning, and ships new receipt-load startup/attribution flags, alongside a large set of breaking changes spanning receipt formats, containment, decoding safety bounds, and exit-code semantics.

└──▷ **WHAT SHIPPED · 21 FEATURES** most completely described first ? what's the number?

01 New receipt-load controls for startup and attribution NEW

90

`receipt-load` gains `--restart-after` to verify receipt continuity across a proxy restart, `--ready-timeout` (or the `RECEIPT_LOAD_VERIFY_TIMEOUT` environment variable) to control how much extra startup time receipt verification allows, and `--destinations per-request` to attribute destination-only receipts to individual planned requests.

Allow extra time for startup verification in slow environments, such as when the registry is behind a high-latency network.

```
$ receipt-load --ready-timeout 120s --restart-after
/var/run/pipelock/restart.sock
```

Attribute destination-only receipts to individual planned requests when auditing per-request coverage in a fleet run.

```
$ receipt-load --destinations per-request --restart-after
/var/run/pipelock/restart.sock
```

— Three named flags with runnable example commands. v3.8.0

02 Expanded encoded-text inspection for DLP and injection scanning

IMPROVED

80

pipelock now decodes decimal character codes, explicit octal and binary codes, quoted-printable escapes, and mixed standard/URL-safe base64 so more encoded text reaches DLP and injection scanning. Long separator-split hex, base64, base64url, base32, and

base32hex values now receive text DLP inspection via a separate shared 8 MiB allowance; exhausting the allowance refuses outbound content.

— Lists decoded formats and the 8 MiB allowance but no tuning flag. v3.8.0

03 **Nearby-keyword gates for custom DLP patterns** NEW 75

Custom DLP patterns can add `keywords` and `keyword_window` fields to require a nearby label within a character window around any match, reducing false positives; built-in patterns do not accept these fields.

Reduce false positives in a custom DLP pattern by requiring a nearby label within a window of characters around any match.

yaml

```yaml
keywords:
  - "api-key"
  - "secret"
keyword_window: 100
```

— Named config fields with an example, but no limit documentation. v3.8.0

04 **Evidence prefix validation reports unproven end** BREAKING 75

A valid open evidence prefix now reports `VALID PREFIX, END UNPROVEN` and exits 3; `--allow-prefix` deliberately accepts a supported prefix without proving its end — update automation that treats signature validity as completeness.

— Exact exit code, message string, and flag given. v3.8.0

05 **New exit codes for verifier install gate and receipt-load** 75

BREAKING

The verifier install gate now exits `3` for a retryable registry access failure and `2` for a package contradiction or unclassified failure; `receipt-load` exits `4` for a setup failure before receipt evidence exists — update scripts that classify these results.

— Exact exit codes and affected commands named. v3.8.0

06 **Metrics and stats endpoints gated by dedicated listener** 70

BREAKING

`/metrics` and `/stats` are only available when `metrics_listen` names a dedicated listener; omitting, nulling, or blanking the value exposes neither route at startup.

— Exact endpoints and config key named. v3.8.0

07 **Default URL length ceiling lowered to 8000 octets** BREAKING

70

The default URL ceiling is now 8000 octets; operators who need the previous limit must set `fetch_proxy.monitoring.max_url_length` explicitly.

— Exact config key and new default value given. v3.8.0

08 **File-only MCP wrappers reported as unconfirmed** BREAKING

70

File-only MCP wrappers are now reported as unconfirmed; `discover`, assessment, and posture output no longer equate a wrapper found in a configuration file with proven runtime protection, and discovery exits nonzero for unconfirmed coverage.

— Names commands and the new exit behavior. v3.8.0

09 **Guard and sandbox honor kernel cross-directory move restriction**

70

BREAKING

Guard and sandbox now honor the kernel's cross-directory move restriction; cross-directory renames and hard links are granted only when the kernel reports Landlock erratum 3 fixed, otherwise refused with `EXDEV`. Older typed Guard proof readers can reject proofs carrying the new field.

— Mechanism and error code named, no config to adjust. v3.8.0

10 **Receipt report format moves to v2** BREAKING

70

Fleet receipt reports now default to v2 — update consumers that assume v1. Signed URL targets in receipts now retain only scheme, host, and port, dropping userinfo, path, query, and fragment. Unicode tag decoding adds receipt normalization views for DLP and injection scanning; older verifier releases reject receipts that use these new views, so verifiers must be upgraded first.

— Three related receipt-format changes named together. v3.8.0

11 **Core DLP floor enforced regardless of custom patterns**

65

BREAKING

`git scan-diff`, `audit`, and the GitHub Action now apply the core DLP floor even with custom patterns or defaults disabled; inline markers and rule suppressions cannot silence core findings.

— Names three surfaces and the enforcement scope. v3.8.0

12 **Decoding-bound refusal for unresolved encoded views** BREAKING

60

Outbound decoding now refuses when decoding depth or work bounds leave a new view uninspected, affecting supported URL, body, header, WebSocket, MCP, and A2A request paths.

— Names affected paths but no config to adjust the bounds. v3.8.0

13 **Agent presets gain named Preset expressions** BREAKING

60

The `claude-code`, `cursor`, `generic-agent`, and `hostile-model` presets now give supplemental expressions `Preset` names — review name-based alerts and suppressions when replacing older YAML.

— All four presets and the new field name are given. v3.8.0

14 **Managed containment units require reinstall** BREAKING

60

Managed containment units need reinstalling after upgrading — run `contain install`, then `contain verify` and `contain doctor`.

— Exact commands to run are given. v3.8.0

THINNER COVERAGE BELOW

15 **Container and Helm images moved to ghcr.io/pipelab-org**

55

BREAKING

Container and Helm chart images now use `ghcr.io/pipelab-org` — update image repositories, registry policies, and Helm OCI URLs.

— Clear migration note with no further mechanism. v3.8.0

16 **Audit mode still enforces by default** BREAKING

55

`mode: audit` with the default `enforce: true` still blocks ordinary URL findings — use the `audit` preset for an observation trial.

— Named config keys and preset, limited explanation. v3.8.0

17 **Expired scanner exceptions lapse instead of blocking startup**

50

BREAKING

Expired temporary scanner exceptions now lapse instead of preventing startup; a valid expired entry emits a warning and scanning resumes for its scope.

— Explains before/after behavior without a config key. v3.8.0

18 **Generated audit configs retain tool-policy blocks** BREAKING

50

Generated audit configurations now retain explicit tool-policy blocks and disable adaptive escalation — review regenerated policy before replacing a custom configuration.

— Describes the change without naming specific config fields. v3.8.0

19 **Codex MCP setup forwards named environment variables** BREAKING

50

Codex MCP setup now forwards environment variables by name — export each required variable in Codex's runtime environment before upgrading.

— Clear migration step but no variable names listed. v3.8.0

20 **Provider credential relief requires supported endpoint and carrier**

50

BREAKING

Provider credential relief now requires a supported endpoint and carrier; a provider-like hostname alone no longer grants relief, and the built-in policy hash changes — refresh pinned hashes when republishing fleet policy.

— Describes the change without listing endpoints or hash values. v3.8.0

21 **Non-finite floating-point settings rejected** BREAKING

45

Non-finite floating-point values (NaN, infinity) in settings are now rejected — replace them with finite values before upgrading.

— Clear rule but no named field or config key. v3.8.0

└──▷ BREAKING ON UPGRADE

! `/metrics` and `/stats` are only available when `metrics_listen` names a dedicated listener; omitting, nulling, or blanking the value exposes neither route at startup.

! Container and Helm chart images now use `ghcr.io/pipelab-org` — update image repositories, registry policies, and Helm OCI URLs.

! Expired temporary scanner exceptions now lapse instead of preventing startup; a valid expired entry emits a warning and scanning resumes for its scope.

- ! Non-finite floating-point values (NaN, infinity) in settings are now rejected — replace them with finite values before upgrading.
- ! The default URL ceiling is now 8000 octets; operators who need the previous limit must set `fetch_proxy.monitoring.max_url_length` explicitly.
- ! Generated audit configurations now retain explicit tool-policy blocks and disable adaptive escalation — review regenerated policy before replacing a custom configuration.
- ! The `claude-code`, `cursor`, `generic-agent`, and `hostile-model` presets now give supplemental expressions `Preset` names — review name-based alerts and suppressions when replacing older YAML.
- ! `git scan-diff`, `audit`, and the GitHub Action now apply the core DLP floor even with custom patterns or defaults disabled; inline markers and rule suppressions cannot silence core findings.
- ! Codex MCP setup now forwards environment variables by name — export each required variable in Codex's runtime environment before upgrading.
- ! Fleet receipt reports now default to v2 — update consumers that assume v1.
- ! File-only MCP wrappers are now reported as unconfirmed; `discover`, assessment, and posture output no longer equate a wrapper found in a configuration file with proven runtime protection, and discovery exits nonzero for unconfirmed coverage.
- ! Managed containment units need reinstalling after upgrading — run `contain install`, then `contain verify` and `contain doctor`.
- ! Guard and sandbox now honor the kernel's cross-directory move restriction; cross-directory renames and hard links are granted only when the kernel reports Landlock erratum 3 fixed, otherwise refused with `EXDEV`. Older typed Guard proof readers can reject proofs carrying the new field.
- ! A valid open evidence prefix now reports `VALID PREFIX, END UNPROVEN` and exits 3; `--allow-prefix` deliberately accepts a supported prefix without proving its end. Update automation that treats signature validity as completeness.
- ! Outbound decoding now refuses when decoding depth or work bounds leave a new view uninspected — supported URL, body, header, WebSocket, MCP, and A2A request paths are affected.
- ! Signed URL targets in receipts now retain only scheme, host, and port, dropping userinfo, path, query, and fragment — update consumers that relied on receipt paths or queries.
- ! `mode: audit` with the default `enforce: true` still blocks ordinary URL findings — use the `audit` preset for an observation trial.
- ! Provider credential relief now requires a supported endpoint and carrier; a provider-like hostname alone no longer grants relief, and the built-in policy hash changes — refresh pinned hashes when republishing fleet policy.
- ! Unicode tag decoding adds receipt normalization views for DLP and injection scanning — upgrade receipt verifiers before consuming the new views, as older verifier releases reject receipts that use them.
- ! The verifier install gate now exits `3` for a retryable registry access failure and `2` for a package contradiction or unclassified failure; `receipt-load` exits `4` for a setup failure before receipt evidence exists — update scripts that classify these results.

WAS THIS USEFUL?  

Certiv

Feedback

Local Model Adoption Is Accelerating. Govern Your Models and Accounts.

NOTES ↗

RANK

Certiv is a complete visibility and control over every AI agent locally installed on the endpoint. Discover, understand, and safeguard every AI action across your organization.

> Certiv now discovers and classifies every AI model running locally on endpoints, lets admins sanction or block any model or provider account org-wide, and extends its runtime security controls to local model traffic.

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

01 **Local model discovery and identification** NEW 82

Certiv now discovers and inventories models running locally on endpoints through runtimes like Ollama, `vLLM`, and LM Studio, surfacing them on the Models page alongside cloud models. Each local model is identified by file contents and labeled Official (matches a known publisher release), Modified (refusal training removed, even if renamed), or Unverified (no match to a known release); abliterated models — open-weight builds with refusal training stripped — are flagged so security teams can see which endpoints run them. The Models page also surfaces fields showing where each model is running, what it actually is, and which endpoints are running it.

Review and block abliterated or unrecognized local models before they are used in an agent workflow — navigate to Models, filter by on-device location, check the Abliterated field, mark offenders Unsanctioned, then set organization-wide enforcement.

📍 In the Certiv console, go to Models › (filter: Running On = Endpoint) › review each model's Modified or Unverified label under the identification fields › mark it Unsanctioned › set the organization-wide policy to Enforce to block unsanctioned models on every endpoint.

— Names identification labels and runtimes but lacks exact UI field names

Local Model Adoption Is Accelerating. G… launch-20261110-52c4791c

02 **Organization-wide enforcement policy** NEW 80

A new organization-wide enforcement policy lets admins choose Enforce (blocks unsanctioned models and accounts on every endpoint) or Audit (records use without blocking), with coverage optionally extending to unreviewed models.

Review every model running across your org — local and cloud — and block any that security has not approved, without writing policies.

📍 In the Certiv console, go to Models page, review each model's lab, attributes (open-weight, modified, abliterated), runtime, and endpoints. Mark each as Sanctioned or Unsanctioned, then set the organization-wide policy to Enforce to block unsanctioned models on every endpoint.

Detect personal provider accounts (e.g. a personal Anthropic or Cursor account) in use on work laptops and block them org-wide.

📍 In the Certiv console, go to the Accounts review view, find accounts flagged as unsanctioned (personal accounts alongside company-issued ones), mark them Unsanctioned, then set the policy to Enforce to block those accounts across the organization.

Review and block abliterated or unrecognized local models before they are used in an agent workflow — navigate to Models, filter by on-device location, check the Abliterated field, mark offenders Unsanctioned, then set organization-wide enforcement.

📍 In the Certiv console, go to Models › (filter: Running On = Endpoint) › review each model's Modified or Unverified label under the identification fields › mark it Unsanctioned › set the organization-wide policy to Enforce to block unsanctioned models on every endpoint.

Catch personal provider accounts (Anthropic, OpenAI, Cursor) used on company laptops and block them without disrupting sanctioned company accounts.

📍 In the Certiv console, go to Accounts › review each account's provider and owner › mark personal accounts Unsanctioned › set the enforcement policy to Enforce to block unsanctioned accounts across the organization.

— Names both policy modes and exact setting location

Local Model Adoption Is Accelerating. G…

---

03 **Sanctioning for models and provider accounts** NEW 72

Administrators can mark every model and account as Sanctioned or Unsanctioned from a single Models page, covering open-weight, proprietary, local, and cloud models. Provider accounts — such as Anthropic, OpenAI, or Cursor — now appear alongside models for review, letting admins detect personal accounts in use on work laptops and mark them Unsanctioned.

Review every model running across your org — local and cloud — and block any that security has not approved, without writing policies.

📍 In the Certiv console, go to Models page, review each model's lab, attributes (open-weight, modified, abliterated), runtime, and endpoints. Mark each as

Sanctioned or Unsanctioned, then set the organization-wide policy to Enforce to block unsanctioned models on every endpoint.

Detect personal provider accounts (e.g. a personal Anthropic or Cursor account) in use on work laptops and block them org-wide.

📍 In the Certiv console, go to the Accounts review view, find accounts flagged as unsanctioned (personal accounts alongside company-issued ones), mark them Unsanctioned, then set the policy to Enforce to block those accounts across the organization.

Catch personal provider accounts (Anthropic, OpenAI, Cursor) used on company laptops and block them without disrupting sanctioned company accounts.

📍 In the Certiv console, go to Accounts › review each account's provider and owner › mark personal accounts Unsanctioned › set the enforcement policy to Enforce to block unsanctioned accounts across the organization.

— Clear UI workflow given but no API or config surface

Local Model Adoption Is Accelerating. G… launch-20261110-52c4791c

THINNER COVERAGE BELOW

04 **Scout runtime security extended to local models** IMPROVED

54

Scout now sits inline between agents and local models, applying the same runtime security — sensitive data controls and tool- and MCP-server-level governance — to local model traffic as it already does to cloud API traffic.

— Describes mechanism but no config or command given

Local Model Adoption Is Accelerating. G…

05 **Local model governance generally available** IMPROVED

13

Local model governance is now available to all Certiv customers.

— Bare availability statement with no further detail

Local Model Adoption Is Accelerating. G…

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

Bots - Partitioned cookie support for Bot Management (2026-11-09) BLOG ↗

RANK

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

> Cloudflare shipped a Rust SDK for feature flag evaluation, faster Stream MP4 downloads, a GitHub repository listing fix for Workers Builds and Pages, and a Bot Management cookie update for cross-site iframe support.

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

01 **Partitioned attribute on Bot Management cookie** IMPROVED 70

Bot Management now adds the `Partitioned` attribute to the `__cf_bm` cookie on compatible HTTPS requests (when `SameSite=None` and `Secure` can also be set), preserving bot-detection continuity for content embedded in cross-site iframes where browsers would otherwise block third-party cookies.

— Names exact cookie attributes and conditions but no config step

Bots - Partitioned cookie support for B…

THINNER COVERAGE BELOW

02 **Flagship Rust SDK for feature flags** NEW 50

Cloudflare adds a Flagship Rust SDK with an async OpenFeature provider supporting HTTP and local evaluation for Rust server applications.

— Names SDK and evaluation modes but no install/usage snippet product docs

03 **User-scoped GitHub repo listing in Workers Builds and Pages** IMPROVED 45

Workers Builds and Pages now list GitHub repositories based on the authenticated user's own GitHub account rather than organization-wide lists.

— Clear before/after but no UI path or config detail product docs

04 **Faster MP4 downloads from Stream** IMPROVED 15

Cloudflare Stream now generates downloadable MP4 files faster without any implementation changes required.

— No mechanism, numbers, or action given product docs

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

Aura

Feedback ⓘ announced · Nov 10, 2026 PRESS ↗

Aura is an all-in-one protection from identity theft and other online scams for the whole family. Works across devices.

Aura adds Solstice, a new AI engine with a household knowledge graph, an 800M-parameter foundation model, and AI agents that act across identity, credit, devices, and scams.

Aura now includes a household knowledge graph that links every person, device, account, and exposure in a family over time, letting the engine connect events months apart into a single threat story rather than three separate alerts.

Aura now runs a proprietary 800-million-parameter foundation model trained on 16 billion de-identified household security events, learning how threats unfold as sequences so household risk surfaces earlier.

NEW INTEGRATION Aura now includes an action plane of AI agents — covering credit and identity, accounts and passwords, devices, children's safety, and scams — that can carry out protection tasks at a member's direction.

Aura now replaces the long setup process with a conversational setup that asks about the household, turns on the most relevant protections first, and tracks what has been done.

Aura now provides a daily briefing and weekly recap summarizing overnight checks, changes, and the weekly state of household protection.

+2 more

NEW DETECTION Specialized ML models for scam, malware, transaction, breach, and child-safety detection

NEW INTEGRATION School-safety signal integration connecting classroom signals to parent view

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 11, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue