

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 17, 2026 · EVERY DAY

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

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

▸ // HOW THIS ISSUE IS MADE

SHOW

VIEW

ISSUE VIEW

scroll



SCROLL

EXPAND
IN
PLACE

PREVIEW
PANE

INFORMATION DENSITY

everything



BRIEFED

EVERYTHING

Every tool in the issue, in full.

all 212 changes

Do you prefer **scroll**?



▸ // CONTENTS SHOW

48 tools

▼ // THE BRIEF HIDE

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

Three unrelated things stand out. ConductorOne put an assistant inside findings remediation and access reviews, so the reviewer no longer

vim ? help

enables context by hand. Veracode's CLI now scans straight from a

Git repo and generates SBOMs. And GitGuardian added 10 new detectors, 2 new analyzers and an SMTP validity checker while cutting false positives in high-volume checkers.

PROTECT

Scan code and fix dependency findings without leaving the repo

Veracode's CLI reads the Git working tree directly and can pull findings back and auto-fix SCA issues in Limited Availability, collapsing the upload-wait-download loop that made AppSec scanning something you ran nightly rather than at commit. Invicti adds one-click bulk retest so a batch of fixed vulns gets confirmed in one action instead of queued one at a time, and flags scans that failed authentication rather than letting them silently report clean.

[Veracode](#) · [Invicti](#) · [ZeroPath](#)

PROTECT

Run an AI coding agent without giving it your laptop

coi drops Claude Code, Codex and others into Incus containers with kernel-level threat monitoring and supply-chain guards, so a prompt-injected agent is confined to a disposable rootfs rather than your SSH keys. SkillSpector binds MCP HTTP listeners to loopback — closing the local-server-exposed-to-the-network hole — and widens pattern detection into shell payloads, binaries and Claude plugin manifests.

[code-on-incus](#) · [SkillSpector](#) · [Cyera](#)

OFFENSIVE SECURITY

SpecterOps SharpHound

Feedback



v2.17.1

NOTES ↗

RANK

C# Data Collector for BloodHound

SharpHound v2.17.1 adds collection of Azure VM IDs from Windows computers during AD enumeration, enabling correlation between on-premises AD objects and Azure infrastructure.

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

01 Azure VM ID collection from Windows computers NEW

45

SharpHound now collects Azure VM IDs from Windows computers during AD enumeration, enabling BloodHound to correlate on-premises AD objects with Azure infrastructure.

– Describes what is collected and its purpose, no mechanism detail. v2.17.1

WAS THIS USEFUL?



nerva

[Feedback](#)[v1.71.0](#)[NOTES ↗](#)[LATEST](#)[STABLE](#)

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

↳ [GET THIS VERSION](#)



```
$ git clone --branch v1.71.0 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.71.0
```

COVERAGE

[NEW](#) detections

[InterSystems IRIS portal fingerprinter](#), [.NET Remoting HTTP channel exposure](#)

[All of nerva's coverage changes →](#)

▼ 1 more release in this issue · 2026-10-15

[v1.70.9](#)[NOTES ↗](#)[STABLE](#)

↳ [GET THIS VERSION](#)



```
$ git clone --branch v1.70.9 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.70.9
```

WAS THIS USEFUL?



Great Scott Gadgets Cynthion

Feedback



sold out · Oct 15, 2026

STORE ↗

A USB test instrument.

- › Cynthion without enclosure (cynthion-board) sold out at shop.hak5.org on 15 Oct 2026.

WAS THIS USEFUL?



xalgorix

Feedback



v4.6.147

NOTES ↗

RANK

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

xalgorix v4.6.147 adds plain-text endpoint list uploads to scanner context, improves session-level token-usage tracking, and trims repeated agent context in long autonomous scans.

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

01 Plain-text endpoint list uploads for scanner context

NEW

37

Scanner context now accepts plain-text endpoint lists in uploaded files, letting practitioners feed raw target lists directly into a scan session.

– Names the capability but no format or size limits given v4.6.147

02 Session-boundary token usage reconciliation

IMPROVED

28

Token usage is now recorded and reconciled at session boundaries, producing more accurate totals across scan sessions.

– Describes behaviour but no metrics, UI, or API surface named v4.6.147

03 Reduced repeated agent context in long scans

IMPROVED

28

Long autonomous scans now carry less repeated agent context while preserving tool schemas and required guidance, lowering noise over extended runs.

– Explains effect but no mechanism or configurable setting named v4.6.147

COVERAGE

NEW

vulnerability check Bounded regex denial-of-service (ReDoS) detection task

All of xalgorix's coverage changes →

v4.6.145

NOTES ↗

STABLE

Route discovery now filters probe artifacts and merges observations for relative and parameterized routes so coverage reflects real application inputs.

↳ GET THIS VERSION



```
$ git clone --branch v4.6.145  
https://github.com/xalgorix/xalgorix.git  
# already have the repo? check out this version:  
$ git checkout v4.6.145
```

- > Route discovery now filters probe artifacts and combines observations for relative and parameterized routes so coverage reflects application inputs.

WAS THIS USEFUL?



APPLICATION SECURITY

ZeroPath

[Feedback](#)[changelog · Oct 17, 2026](#)[CHANGELOG ▸](#)[RANK](#)

ZeroPath is an AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives.

ZeroPath added new issue-search fields, a `scans.list` endpoint, and self-serve SBOM and audit-log exports, while tightening permissions around re-investigation, tag management, and programmatic API key access in ways that will break existing integrations.

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

01 Organization-wide SBOM export in CycloneDX format NEW

88

Exports the organisation's full SBOM inventory via Reporting > Organization SBOM, selecting CycloneDX 1.6 JSON, which produces a ZIP containing one SBOM per repository plus a manifest listing included and skipped repositories.

Download your organisation's full SBOM inventory in CycloneDX format after a release, to feed into a downstream SCA pipeline.

📌 In the app, go to Reporting > Organization SBOM, select CycloneDX 1.6 JSON, then click Export. The downloaded ZIP contains one SBOM per repository plus a manifest listing included and skipped repositories.

– Exact UI path, format version and output structure given [changelog-20261017-e49b8332](#)

02 Issue search filters and sorting IMPROVED

60

Issue search now accepts a `presentation: true` field and a `knownVulnerableVersion` field, and can filter results by lane and sort by priority.

– Named API fields but no endpoint path or example [changelog-20261017-e49b8332](#)

THINNER COVERAGE BELOW

03 `scans.list` endpoint with dependency findings counts NEW

45

Adds `scans.list`, which returns issue counts that now include dependency findings.

– New endpoint named but usage not detailed [changelog-20261017-e49b8332](#)

↳ BREAKING ON UPGRADE

! Re-investigation from the app, API, MCP, or Assistant now requires [Close Issues](#) permission on the finding's repository; members who previously could re-investigate

without it will receive a 403 until the permission is granted.

- ! Creating, editing, and deleting repository tags now requires [Manage Teams](#) or [Edit SAST Configuration](#) ; members who previously had tag management access without these permissions will lose it.
- ! API key create, list, and delete operations now return 403 for requests authenticated with an API key or [Authorization](#) header; callers that managed keys programmatically must switch to a signed-in browser session.
- ! Tag create/update requests that include the (unspecified) previously accepted field are now refused; callers must use the explicit include, exclude, and release change fields instead.

WAS THIS USEFUL?



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

GitHub shipped general availability of open-source license compliance for Enterprise Cloud, giving organizations policy enforcement, new APIs, and a dedicated reviewer role, alongside multi-repository support and context improvements for the GitHub app in Slack and Teams.

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

01 Open source license compliance (GA) NEW 90

GitHub Enterprise Cloud's open-source license compliance is now generally available. It adds new API endpoints for downloading, updating, and replacing the enterprise license policy, plus a separate endpoint for checking whether individual packages comply with policy. A new **Require license compliance results before merging** ruleset condition enforces the enterprise-wide license policy on pull requests that add or modify dependencies. A new predefined role, **Enterprise Open-Source License Manager**, is available for reviewing and approving license closure requests, and license policy changes now emit audit log events for external systems of record. The dependency graph and Software Bill of Materials (SBOM) exports are now populated with license data for GitHub Enterprise Cloud data-residency customers, and coding agents can use the new package compliance check endpoint to avoid recommending packages that violate an organization's license policy.

– Names endpoints, ruleset condition, and role, but no exact HTTP methods/paths given.

Open source license compliance generall...

02 Multi-repository support in GitHub for Slack and Teams

IMPROVED

65

Users can now work across multiple repositories in a single @GitHub conversation in Slack or Teams, creating issues, opening pull requests, and investigating bugs without switching context. The integration also supports shared Slack messages and threads as context for @GitHub, reducing the need to copy and paste discussion details, and repository switching mid-conversation has been improved.

▶ Watch video

Triage a bug across frontend and backend repos in one Slack thread without opening GitHub.

📌 Steps: 1. In any Slack thread, type: '@GitHub Investigate this issue using the code in our frontend and backend repositories.' 2. @GitHub will pull context from both repositories and the current thread to surface relevant code and findings inline.

Get a cross-repo PR review queue from Slack to prioritise your review work each morning.

📌 Steps: 1. In any Slack channel or DM, type: '@GitHub Show me all my open PRs across repositories and highlight which ones need my attention.' 2. @GitHub responds with a consolidated list grouped by urgency, without leaving Slack.

– Concrete example commands given, but underlying mechanism of repo switching unspecified.

Multi-repository support and other impr...

COVERAGE

NEW supported platform [Xcode 27 runner image \(GA\)](#).

[All of GitHub's coverage changes →](#)

Xcode 27 runner image generally available (2026-10-16)

BLOG ↗

STABLE

GitHub Actions introduces a new Xcode-based image model where each image targets one major Xcode version via `runs-on:` labels rather than an OS version.

- › Adds `xcode-27` and `xcode-27-xlarge` as `runs-on:` labels for GitHub-hosted Apple Silicon (arm64) runners shipping macOS 27 with the Xcode 27 toolchain and Apple SDKs.

WAS THIS USEFUL?



skylos

[Feedback](#)[v4.50.0](#)[NOTES ↗](#)[RANK](#)

Open source local-first PR scanner that finds dead code, security bugs, secrets, quality regressions, and AI-code mistakes before merge. For first timers refer to <https://duriantaco.github.io/skylos/repo-map/>

Skylos this window bundled native Go engines directly into its PyPI wheels, container image and GitHub Action, added fine-grained upload/telemetry consent controls, and significantly hardened the `skylos done` agent gate with human approval, AI-code gating, and broader test coverage checks, alongside a new `skylos cisd init` command and clearer scan output.

←→ **WHAT SHIPPED** • **10 FEATURES** most completely described first ? what's the number?

01 Native Go engine bundling across platforms **NEW** 83

Skylos now bundles native Go dead-code and security engines directly in PyPI wheels for Linux, macOS, and Windows (x86_64 and arm64), the container image, and the GitHub Action, eliminating the need to build from source on supported platforms. The `SKYLOS_GO_BIN` environment variable lets Skylos use a separately built Go engine on unsupported platforms, and `skylos doctor` now prints the build command needed for the native Go engine where the bundled wheel is unavailable.

Point Skylos at a manually compiled Go engine on an unsupported platform (e.g. Alpine Linux / musl).

```
$ SKYLOS_GO_BIN=/usr/local/bin/skylos-go skylos . -a
```

— Names platforms, env var and doctor output with a runnable example [v4.50.0](#)

02 Agent done-gate and verification enhancements **IMPROVED** 80

The agent gate subsystem gained several behaviors across releases: guardrail and AI-agent workflow rules enforcing deployment and behavior policies on agent-generated branches; the `done` gate now fails by default on critical AI-code issues and always prints the verdict; `skylos agent` now reports silenced checks and the `done` gate no longer falsely blocks because of them; the gate blocks code that special-cases tests (e.g. bypassing checks only while tests run); `skylos done` was extended to check JavaScript and TypeScript tests beyond Python; and `skylos agent verify` gained session-stop verification and changed-line test checks.

Check whether an agent-authored branch is ready to merge — the gate now fails automatically on any critical AI-code issue and always prints the verdict.

```
$ skylos done --base main
```

– Enumerates every gate behavior and command with a runnable example v4.44.0 v4.45.0
v4.46.0 v4.47.0 v4.48.0 v4.49.0

03 Upload consent and trust controls NEW 78

The new `--no-upload` flag suppresses telemetry uploads, takes precedence over all other upload settings, and omits the OS username from run summaries. `skylos feedback` now shows an explicit agent-run upload consent prompt, and the `cloud` command's new `--require-trusted-upload` flag gates uploads on trust status, with upload trust now shown in its output.

Run a full scan in CI without sending any telemetry or leaking the OS username in run summaries.

```
$ skylos . -a --no-upload
```

– Three named flags/commands with a runnable example for one v4.44.0 v4.50.0

04 `skylos cicd init` CI/CD workflow generation NEW 60

The new `skylos cicd init` subcommand generates a working gated-PR CI/CD workflow, and it also prints completion settings for Vitest and Jest as part of its output.

– Names the exact subcommand but no sample invocation shown v4.46.0 v4.48.0

THINNER COVERAGE BELOW

05 Confidence and severity threshold adjustments IMPROVED 56

Unused-parameter and experimental Rust findings below the default confidence threshold are now hidden, and Go rules `G203` and `G260` were adjusted to MEDIUM severity while `G206` was set to LOW.

– Names exact rule IDs and severities but no action to take v4.50.0

06 Scan output clarity and assessment changes IMPROVED 55

The CLI now explains security evidence and scan coverage directly in its output, replaces blended A-F grades with clear scan assessments for source scans, and scans now skip generated files while reserving critical severity for shipped code.

Run a full scan to see the new security evidence explanations and scan coverage detail alongside findings.

```
$ skylos . -a
```

– Names concrete output changes with a runnable example command v4.48.0

07

'skylos done' risk gating and tracking enhancements NEW

50

The `skylos done` agent workflow now requires a human to unblock the risk-review check before a risky PR can be merged, and it now tracks failure origin and detects orphaned work, surfacing which changes caused a done-check failure and identifying orphaned code.

– Describes gating mechanism but no example or command shown Names the command and behavior but no usage example

v4.49.0

08

PR review workflow improvements IMPROVED

38

Skylos now supports merge queues in PR review workflows, reuses PR review summaries, and avoids posting repeated inline comments.

– Names behaviors but gives no mechanism or example v4.50.0

09

Secret-printing shell blockers NEW

28

Adds hooks that block shell commands which print secrets.

– Bare description of a hook with no mechanism or example v4.48.0

10

Provenance tracking for agent edits NEW

25

Records agent edits and preserves commit associations in provenance tracking.

– Thin one-line description with no mechanism or usage v4.48.0

WAS THIS USEFUL?

Veracode

[Feedback](#)[CLI updates \(2026-10-13\)](#)[BLOG](#)[RANK](#)

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

Veracode CLI v2.53.0 adds Git-native static scanning and findings retrieval, a Limited Availability SCA auto-fix command, SBOM generation, and a long list of packaging and policy enhancements, alongside a breaking change to the default SCA download URL.

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

01 Static scan and cached findings commands **NEW** 80

Introduces `veracode scan static` to run static scans and review security findings within a Git workflow, and `veracode findings static` to access cached static scan findings from the local Veracode directory for fast CI feedback without re-scanning.

Run a static scan within your Git workflow and then retrieve cached findings from the local Veracode directory for fast CI feedback without re-scanning.

```
$ veracode scan static && veracode findings static
```

— Two named commands with a runnable example [CLI updates \(2026-10-13\)](#)

02 Packaging enhancements in veracode package **IMPROVED** 80

Adds to `veracode package`: `--csharp-source` to package C# source directly for Static Analysis without compiling to binaries; generic packager support for Apex; `--skip-packagers` to exclude specific artifacts without affecting other languages; support for `requirements-dev.txt` in Python workflows; exit codes 5, 6, and 7 for CI/CD automation; `--exclude` to omit Java packages or JAR files from scanning; the `VERACODE_PACKAGE_JAVA_SOURCE=true` environment variable and `--java-source` flag for Java source packaging; the `VERACODE_PACKAGE_GOLANG_GENERATE` environment variable for client-side CGO generation in Go apps; the `SRCLR_MAKE_WINDOWS` environment variable to preprocess C/C++ with Make on Windows using native gcc or clang; automatic inclusion of `.yaml`, `.yml`, and `.properties` files when packaging Java source; support for the Poetry package manager for Python autopackaging; support for the MSBuild .NET framework; and support for Go v1.24.

— Many named flags and env vars, no example shown [CLI updates \(2026-10-13\)](#)

03 New SCA commands for fixes and SBOM generation **NEW**

Introduces `veracode fix sca` (Limited Availability) to generate fixes for vulnerable dependencies identified in SCA scans, and adds SBOM generation support to `veracode sca` via SCA agent-based scans, with output in CycloneDX 1.6 and SPDX 2.3 formats.

Generate SCA-based fixes for vulnerable dependencies in a JavaScript project during a CI pipeline run.

```
$ veracode fix sca
```

– Limited availability command shown with example usage Named formats and command, no example shown

CLI updates (2026-10-13)

04 Policy, vulnerability database and parsing updates in veracode scan

IMPROVED

65

Enhances `veracode scan` with support for a custom policy defined in the Veracode Platform and viewing policy evaluation for individual findings alongside the overall policy result, specifying a custom location from which vulnerability databases can be downloaded, parsing Dockerfile heredoc syntaxes, categorizing secret-type findings as Critical in scan results, and prioritizing vulnerability results from the NVD database when available.

– Several named behaviours, no example command given

CLI updates (2026-10-13)

05 Default SCA download URL changed

BREAKING

65

The `SRCCLR_DOWNLOAD_URL` environment variable default is now `https://sca-downloads.veracode.com`; the previous URL `https://download.sourceclear.com` is deprecated and all instances should be updated.

– Exact env var and URLs with clear migration step

CLI updates (2026-10-13)

THINNER COVERAGE BELOW

06 Toolchain discovery via veracode package discover

NEW

55

Introduces `veracode package discover` to detect toolchains required to create scannable artifacts, later enhanced with toolchain detection for JavaScript projects using npm or Yarn, toolchain detection for Java projects using the Gradle package manager, and new exit codes.

– Command named, no example or exit code values

CLI updates (2026-10-13)

07 Target UUID support for dynamic scans IMPROVED 45

`veracode dynamic scan` now supports target UUIDs for triggering scans, while webhook UUIDs continue to work.

– Short description, no example given CLI updates (2026-10-13)

08 OAuth2 login onboarding command NEW 40

Introduces `auth login` command to onboard new users via OAuth2.

– Bare command description, no example or scope detail CLI updates (2026-10-13)

↳ **BREAKING ON UPGRADE**

! The `SRCCLR_DOWNLOAD_URL` environment variable default is now `https://sca-downloads.veracode.com`; the previous URL `https://download.sourceclear.com` is deprecated and all instances should be updated.

WAS THIS USEFUL?



Invicti

Feedback



changelog · Oct 17, 2026

CHANGELOG ↗

RANK

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

Invicti's Platform On-Demand update adds API key authentication for CI/CD Source Scans, one-click bulk vulnerability retest, gRPC scan support, authentication-failure detection for scans, clearer API catalog origin tracking, and SCIM bulk user removal.

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

01 Authentication-failure scan status NEW 60

Adds a scan failure mode triggered by authentication problems, so scans that never ran authenticated are flagged rather than silently returning incomplete results; the option can be enabled in scan policy or scan settings to fail a scan outright on authentication issues.

Fail a pipeline stage immediately when a scan runs without successful authentication, preventing false-clean results.

📌 In the scan policy or scan settings, enable the option to fail a scan on authentication issues so the scan reports failure instead of returning unauthenticated results.

– Describes mechanism and where to enable it, but not the exact setting name.

changelog-20261017-8e05cd1b

02 gRPC scan compatibility NEW 60

Adds gRPC compatibility by sending the trailers header and an empty gRPC frame for parameterless RPCs.

– Names exact protocol mechanism but gives no setup steps. changelog-20261017-8e05cd1b

THINNER COVERAGE BELOW

03 API catalog discovery-source visibility NEW 50

Adds target name and URL display and filtering, plus per-source discovery icons and a discovery-source column, to the API catalog so each API's origin is immediately visible.

– Names concrete UI additions but no navigation path. changelog-20261017-8e05cd1b

↳ ALSO FROM THESE RELEASES

:

.

WAS THIS USEFUL?



fcaptcha

[Feedback](#)

v1.46.0

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

Self-hosted, invisible CAPTCHA that detects AI agents driving real browsers. 2 AA, and EAA: no cookies, no third-party data sharing, and no visual or audio puzzle.

fcaptcha shipped a Go server-side client, drop-in script-compatible endpoints for Turnstile/hCaptcha/reCAPTCHA, automatic hidden-field token injection, a new passive pointer-movement observation, and a breaking change to how scores are reported on the reCAPTCHA-compatible verify endpoint.

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

01 Drop-in API paths for Turnstile, hCaptcha, reCAPTCHA

NEW

85

The widget is now also served at `/turnstile/v0/api.js`, `/1/api.js`, and `/recaptcha/api.js`, each exposing that provider's JS API (`render`, `getResponse`, `reset`, `remove`, `execute`, `ready`) and matching form-field behaviour, enabling drop-in replacement of Turnstile, hCaptcha, or reCAPTCHA by swapping only the script URL and siteverify endpoint.

Drop FCaptcha into an app built for Cloudflare Turnstile by pointing only the script src and siteverify URL at your FCaptcha instance — no other code changes needed.

 text

```
<script src="https://<your-fcaptcha-host>/turnstile/v0/api.js"
  async defer></script>
```

– Lists exact endpoints, JS API methods and swap mechanism

v1.45.0

02 Go client package for server-side verification NEW

83

Adds a Go client package (`go get github.com/WebDecoy/FCaptcha`) with `Client.Verify` to redeem a CAPTCHA token and `Client.Middleware` to guard an `http.Handler` , supporting optional action, hostname, and score checks. Uses only the Go standard library and fails closed.

Install the Go client package to start verifying fcaptcha tokens server-side.

```
$ go get github.com/WebDecoy/FCaptcha
```

– Names exact package, methods and install command v1.43.0

03 reCAPTCHA-compatible score semantics inverted

BREAKING

83

On the `/recaptcha/api/siteverify` path, `score` is now inverted to follow reCAPTCHA v3 convention (higher = more human, `1 - score`) so existing `score >= threshold` checks work; FCaptcha's original score is still returned as `fcaptcha_score` . Turnstile and hCaptcha paths are unchanged, but any caller comparing the old FCaptcha-native score against a threshold must re-evaluate it.

Verify a token via the reCAPTCHA-compatible endpoint and read both the normalised `score` and the raw `fcaptcha_score` for side-by-side comparison.

```
$ curl -s -X POST https://<your-fcaptcha-host>/recaptcha/api/siteverify \
  -d 'secret=<secret>&response=<token>' | jq '{score, fcaptcha_score}'
```

– Exact endpoint, field names and migration note given v1.46.0

04 Automatic hidden-field token injection in forms NEW

78

A checkbox widget inside a `<form>` now automatically injects a hidden `fcaptcha_token` input, filled on success and cleared on reset or expiry. Rename it with `responseFieldName` / `data-response-field-name` , or disable it with `responseField: false` / `data-response-field="false"` .

– Names config/data attributes but lacks a runnable example v1.45.0

05 Pointer-movement observation signal NEW

75

Adds `pointer-movement-v1` to `experimental.observations`, reporting whether pointer moves ever show zero movement as real mice and trackpads do; this is observe-only and does not affect the score or token decision.

Enable the new pointer-movement observation in your fcaptcha config to gather passive mouse/trackpad data without affecting scoring.

```
experimental:
  observations:
    - pointer-movement-v1
```

– Config key and behaviour named, but explicitly non-scoring v1.46.0

↳ **BREAKING ON UPGRADE**

! On the `/recaptcha/api/siteverify` path, `score` is now inverted to match reCAPTCHA v3 semantics (higher = more human). Any caller that compared the old FCaptcha-native score against a threshold must re-evaluate that threshold; the original score is still available as `fcaptcha_score`.

WAS THIS USEFUL?



StepSecurity

Feedback



Secure Registry: Packagist (PHP) Support (Beta) (2026-10-16)

BLOG ↗

StepSecurity is a security platform that hardens CI/CD pipelines, such as GitHub Actions, by detecting supply-chain risks, enforcing least-privilege runners, and monitoring workflow tampering.

Packagist Registry Settings panel adds Dist URL Rewriting, Keep Git Source, Block Branch Versions, and `::dev::gha:` API-key source attribution for Composer.

- › Adds **Dist URL Rewriting** setting to route package zip downloads through Secure Registry so every archive is checked against your policy.
- › Adds **Keep Git Source** setting to preserve git clone entries in Composer package metadata.
- › Adds **Block Branch Versions** setting to hide branch versions from Composer package metadata.
- › Adds source attribution for Composer: appending `::dev:` or `::gha:` to the API key in `http-basic` or `COMPOSER_AUTH` traces requests to a developer machine or GitHub Actions run in Policy Evaluations.

COVERAGE

NEW

integration [Packagist \(PHP\) Secure Registry support](#)

[All of StepSecurity's coverage changes →](#)

WAS THIS USEFUL?



GitGuardian Detection Engine

[Feedback](#)[Detection Engine Updates Version 2.174 \(2026-10-13\)](#)[BLOG](#)[RANK](#)

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

GitGuardian's Detection Engine 2.174 expands coverage with 10 new detectors, 2 new analyzers and an SMTP validity checker, broadens detection scope and custom-host support across dozens of existing checkers, cuts false positives in several high-volume detectors, and consolidates or removes a handful of detectors in a breaking cleanup.

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

01 Detector consolidation and removals

BREAKING

85

`azure_event_grid_access_key_with_host` merged into `azure_event_grid_access_key`, `grafana_service_account_token_with_host` merged into `grafana_service_account_token`, `kubernetes_jwt_with_host` merged into `kubernetes_jwt`, `planetscale_database_password_with_host_and_username` merged into `planetscale_database_password`, and `postgres_assignment_attached_port` / `postgres_assignment_no_port` merged into `postgres_assignment`; the `line_notify_token` detector and checker were removed entirely since the LINE Notify service is no longer available; and `ftp_credentials_assignment` now captures only FTP secrets, with SMTP handled by the new `sftp_credentials_assignment` detector. References to the removed detector names will no longer resolve.

– Exact removed/merged names given, clear migration note for integrators

[Detection Engine Updates Version 2.174 ...](#)

02 Extended detection scope for existing detectors

IMPROVED

75

`azure_entra_app_secret` detector now detects credentials declared against Microsoft national cloud sign-in hosts; GitHub Server-to-server token detection (`ghs_<app_id>_<JWT>`) now covers any number of scopes including longer tokens previously missed; `docker_credentials` detector supports the new Docker credential format; `ssh_password` detector detects passwords sent in expect scripts after an SSH password prompt; `mssql_assignment` detector detects MSSQL credentials assigned without a port; `ftp_credentials_assignment` checker handles

secrets without a port; and Gemfury detectors detect credentials outside Gemfury repository URLs.

– Before/after behaviour named per detector, no user-facing command

Detection Engine Updates Version 2.174 ...

03 Custom host and proxy support across checkers **IMPROVED**

75

`azure_entra_access_token` and Azure Active Directory API Keys checkers now accept a custom host; Azure Entra ID token, GitLab Deploy Token, and LiteLLM API key checkers can reach custom hosts through a proxy; and `n8n` API key and MCP access token checkers now validate keys and tokens found with a host written without its scheme.

– Names every checker gaining custom-host or proxy support

Detection Engine Updates Version 2.174 ...

04 False positive reduction across multiple detectors **IMPROVED**

75

`generic_high_entropy_secret` no longer reports image names from known container registries; `generic_database_assignment` no longer flags shell command substitutions, empty defaults, the `--usersn` flag, or malformed host or username values; and `microsoft_azure_storage_account_key` no longer flags JFrog reference tokens assigned to `jfrog` variables.

– Specific causes of false positives named per detector, no runnable action

Detection Engine Updates Version 2.174 ...

05 New detectors, analyzers and SMTP validity checker **NEW**

70

Adds new detectors for `sftp_credentials_assignment`, Lago API keys, Composio project API keys, Composio organization API keys, Codex API keys, CARTO Basemaps API keys, Google Chat incoming webhook URLs, HubSpot CLI personal access keys, AmneziaWG VPN private keys, and Kubernetes bearer tokens stored in kubeconfig files; adds new analyzers for GoCardless API access tokens and Adobe refresh tokens; and enables a validity checker for SMTP credentials.

– Names every new detector but no usage mechanism beyond detection itself

Detection Engine Updates Version 2.174 ...

06 Expanded region and format support in existing checkers

IMPROVED

70

`snyc_key_v2` checker now supports the SNYK-US-02 region,
`snowpark_api_credentials` checker now supports short account identifiers,

hostnames, and full Snowflake URLs, and `algolia_monitoring` checker now validates all Algolia keys with more permissions than the default search key.

– Specific checkers and scope named, no mechanism beyond coverage change

Detection Engine Updates Version 2.174 ...

↳ **BREAKING ON UPGRADE**

- ! The `line_notify_token` detector and checker have been removed (the LINE Notify service is no longer available).
- ! The `azure_event_grid_access_key_with_host` detector is removed — merged into `azure_event_grid_access_key`; any references to the old detector name will no longer resolve.
- ! The `grafana_service_account_token_with_host` detector is removed — merged into `grafana_service_account_token`.
- ! The `kubernetes_jwt_with_host` detector is removed — merged into `kubernetes_jwt`.
- ! The `planetscale_database_password_with_host_and_username` detector is removed — merged into `planetscale_database_password`.
- ! The `postgres_assignment_attached_port` and `postgres_assignment_no_port` detectors are removed — merged into `postgres_assignment`.
- ! The `ftp_credentials_assignment` detector now captures only FTP secrets, not SFTP; SFTP is handled by the new `sftp_credentials_assignment` detector.

WAS THIS USEFUL?



REVERSE ENGINEERING, MALWARE & FORENSICS

Mandiant capa

[Feedback](#)[v10.0.0](#)[NOTES ↗](#)[CODE ↗](#)

The FLARE team's open-source tool to identify capabilities in executable files.

Adds `--reproducible` freeze flag for deterministic output and gates vivisect workspace loading behind `CAPA_LOAD_VIV_WORKSPACE=1`.

↳ GET THIS VERSION

```
$ git clone --branch v10.0.0 https://github.com/mandiant/capa.git
# already have the repo? check out this version:
$ git checkout v10.0.0
```

- › Adds `--reproducible` flag to the `freeze` subcommand that zeros dynamic header metadata for deterministic output.
- › Adds `CAPA_LOAD_VIV_WORKSPACE=1` environment variable to opt in to loading existing `<sample>.viv` vivisect workspace files, which are no longer loaded by default due to pickle deserialization risk.

↳ BREAKING ON UPGRADE

- ! The `os_` parameter is removed from `capa.loader.collect_metadata()` — external scripts passing this argument must be updated; OS metadata is now sourced exclusively from `extract_global_features()`.
- ! Existing `<sample>.viv` vivisect workspace files are no longer loaded automatically; set `CAPA_LOAD_VIV_WORKSPACE=1` to restore the previous behaviour.

COVERAGE

NEW detection [12 new capa detection rules](#)

[All of capa's coverage changes →](#)

WAS THIS USEFUL?



Mandiant FLOSS

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

FLARE Obfuscated String Solver - Automatically extract obfuscated strings from malware.

Adds `FLOSS_LOAD_VIV_WORKSPACE=1` env var for safe opt-in vivisect workspace loading; migrates Ghidra import script to PyGhidra.

↳ GET THIS VERSION

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

- › Adds `FLOSS_LOAD_VIV_WORKSPACE=1` environment variable to explicitly opt in to loading `<sample>.viv` workspace files, replacing the previous unsafe auto-load behaviour.
- › Migrates the Ghidra import script to PyGhidra.

↳ BREAKING ON UPGRADE

- ! Python 3.10 or newer is now required; earlier Python versions are no longer supported.
- ! `FLOSS_SAVE_WORKSPACE` now only accepts the values `1`, `true`, `yes`, or `on`; other truthy values that previously worked will break.
- ! FLOSS no longer auto-loads `<sample>.viv` workspace files by default; existing workflows that relied on automatic `.viv` loading must now set `FLOSS_LOAD_VIV_WORKSPACE=1` explicitly.

COVERAGE

NEW supported platform [Basic ELF binary support](#)

[All of FLOSS's coverage changes →](#)

WAS THIS USEFUL?



Finite State Platform

[Feedback](#)

October 14, 2026 - Version counts, component pages, Policy Overview, and re-scan (2026-1

[BLOG](#)[RANK](#)

Finite State Platform is a firmware and software supply chain security tool that analyzes binaries and embedded devices to identify vulnerabilities, SBOM components, and risks.

Finite State added API-level signaling to distinguish busy version comparisons and refused re-scans from genuine failures, scaled component list pages to 10,000 results, and made Policy Overview rows clickable for faster navigation to findings.

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

01 Retry-After signaling for busy version comparisons

NEW

90

Adds `503 COMPARISON_BUSY` with a `Retry-After: 2` header from `GET /api/public/v0/versions/{baseVersionId}/compare/{targetVersionId}` when too many Version comparisons are running at once, enabling callers to implement retry logic instead of treating the throttle as a failure.

Implement retry logic in a polling loop for version comparisons that may be throttled under heavy load.

```
$ curl -i -X GET
'https://platform.finitestate.io/api/public/v0/versions/{baseVersionId}/compare/{targetVersionId}' \
-H 'Authorization: Bearer <token>' \
# On 503 COMPARISON_BUSY, read Retry-After: 2 and sleep before
retrying
```

– Full endpoint, status code, retry header, and example given

October 14, 2026 - Version counts, comp...

02 Conflict signaling for refused re-scans

NEW

90

Adds `409 SCAN_CONFIG_CONFLICT` from `POST /api/public/v0/versions/{projectVersionId}/rescan` when a basic scan and Quick scan are combined before a Version is renamed, making a refused re-scan distinguishable from one that actually started.

Detect a refused re-scan before a Version is renamed by catching the 409

SCAN_CONFIG_CONFLICT response.

```
$ curl -i -X POST
'https://platform.finitestate.io/api/public/v0/versions/{projectId}/rescan' \
-H 'Authorization: Bearer <token>' \
-H 'Content-Type: application/json'
# HTTP 409 SCAN_CONFIG_CONFLICT means the re-scan was refused –
no new Scan started
```

– Endpoint, status code, trigger condition, and example given

October 14, 2026 - Version counts, comp...

THINNER COVERAGE BELOW

03 Larger component list pages **IMPROVED**

55

GET /api/public/v0/components now loads a page of 10,000 components within the request time limit.

– Endpoint and page size named but no example given

October 14, 2026 - Version counts, comp...

04 Clickable rows in Policy Overview **IMPROVED**

50

Policy Overview severity and exploit rows now open the matching component or finding list when clicked anywhere on the row, and License Violations/Warnings rows open only the findings for that specific license policy.

– UI behavior described without an exact navigation path

October 14, 2026 - Version counts, comp...

WAS THIS USEFUL?



witr

Feedback



v0.3.4

NOTES ↗

Why is this running? Trace any process, port, container, or file back to what started it - CLI + TUI.

witr v0.3.4 adds an interactive browser playground with guided tutorial and anonymous analytics via GoatCounter.

↳ GET THIS VERSION

```
$ git clone --branch v0.3.4 https://github.com/pranshuparmar/witr.git
# already have the repo? check out this version:
$ git checkout v0.3.4
```

- › Adds an interactive browser playground with a guided tutorial at <https://pranshuparmar.github.io/witr/> — no install required.
- › Adds anonymous analytics via GoatCounter.

COVERAGE

NEW

supported platforms

[Apple container runtime support](#), [LoongArch \(loong64\) architecture support](#)

[All of witr's coverage changes](#) →

WAS THIS USEFUL?



NETWORK SECURITY

Sophos Firewall

Feedback



Version 22.0 MR3Build 576Released on October 16, 2026

NOTES

RANK

Sophos Firewall is a next-generation network security appliance and software that protects networks through firewalling, intrusion prevention, VPN, web filtering, and threat detection.

Sophos Firewall 22.0 MR3 adds a CLI option to relax strict IPsec SA policy enforcement for third-party VPN peers, alongside two upgrade and hardware-support breaking changes.

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

01 CLI command to disable strict IPsec SA checks NEW

63

A new CLI command turns off strict policy checks for site-to-site IPsec VPN tunnels when third-party peers share Security Associations (SAs) across multiple subnet pairs, preventing incoming IPsec traffic from being dropped.

– Names mechanism and symptom fixed but not exact command syntax

Version 22.0 MR3Build 576Released on Oc...

THINNER COVERAGE BELOW

02 Upgrade path from 22.0 MR3 to 23.0 EAP1 blocked

BREAKING

40

Upgrading to 23.0 EAP1 from version 22.0 MR3 Build 576 is not supported.

– Clear constraint but no workaround or reasoning given

Version 22.0 MR3Build 576Released on Oc...

03 XG and SG hardware appliances dropped

BREAKING

35

SFOS 22.0 GA and later versions no longer support XG and SG Series hardware appliances.

– States the cutoff but no migration guidance

Version 22.0 MR3Build 576Released on Oc...

BREAKING ON UPGRADE

! Upgrading to 23.0 EAP1 from version 22.0 MR3 Build 576 is not supported.

! SFOS 22.0 GA and later versions no longer support XG and SG Series hardware appliances.

WAS THIS USEFUL?



◆ IDS / IPS

maltrail

Feedback



3.5

NOTES

RANK

Malicious traffic detection system

Maltrail 3.5 focuses on running lean on constrained hardware, slashing sensor memory use and adding ARMv6, Windows and NetBSD/OpenBSD support, alongside a small dashboard improvement.

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

01 Sensor memory footprint reduction IMPROVED 85

Streams the static trail set from its cache instead of loading it into memory whole, cutting trail-update peak RAM from 1,144 MB to 69 MB and sensor steady-state from 153 MB to 32 MB, enabling operation on a 256 MB Raspberry Pi. On hosts with under 768 MB available (including container and systemd limits), the trail merge spills to a SQLite file while keeping the same byte-identical `trails.csv` output. `UPDATE_SERVER` downloads are now streamed to disk rather than buffered in memory.

– Concrete before/after numbers and named mechanisms, but no manual flag to invoke 3.5

02 Expanded sensor platform support NEW 80

Adds a static ARMv6 binary (`arm-unknown-linux-musleabihf`) so the sensor runs on every 32-bit Raspberry Pi — Pi 1 and Zero included — under Raspberry Pi OS, DietPi, or Alpine, including on a 64-bit kernel with a 32-bit userland; `install.sh` selects it automatically. Adds an `install.ps1` Windows installer for the sensor, and adds NetBSD and OpenBSD to the tested platform matrix.

– Names exact binary, installer scripts and target OSes 3.5

THINNER COVERAGE BELOW

03 Sensor status display on dashboard NEW 20

Adds a basic sensor status display on the dashboard.

– No detail on what the status shows or where 3.5

WAS THIS USEFUL?



CLOUD & CONTAINER SECURITY

aya

[Feedback](#)[aya-build-v0.2.1](#)[NOTES](#)[RANK](#)

Aya is an eBPF library for the Rust programming language, built with a focus on developer experience and operability.

Aya's aya-build v0.2.1 release adds multi-uprobe attach support, a programmatic `RingBuf::create()` API, perf counter reading via a decoupled `PerfEventArray`, ksym program loading, and new attachment methods for TC, CgroupSockAddr and sk_skb/sk_msg programs, while removing Borrow-based generics from map APIs as a breaking change.

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

01 Attachment API additions for TC, CgroupSockAddr, and sk_skb/sk_msg

NEW

75

Adds a `set_tc_classid` method to `TcContext` and `SkBuff` for setting TC class IDs from eBPF programs, adds `attach_to_link` for `CgroupSockAddr` programs, and uses `FdLink` for `sk_skb` and `sk_msg` program attachment.

– Names every method and type added across targets [aya-build-v0.2.1](#)

02 Map methods drop Borrow-based generic APIs

BREAKING

70

Map methods in `aya-ebpf` and userspace no longer accept types via Borrow traits — callers that relied on Borrow-based generics on map get/insert/remove methods must pass the concrete key/value types directly.

– Clear migration note naming affected methods and crates [aya-build-v0.2.1](#)

03 Multi-uprobe attach support with kernel feature probe

NEW

60

Adds multi-uprobe attach support, enabling a single uprobe program to attach to multiple targets, along with a kernel feature probe to detect multi-uprobe support on the running kernel.

– Mechanism described but no exact API call named [aya-build-v0.2.1](#)

04 RingBuf::create() for programmatic ring buffer maps

NEW

60

Adds `RingBuf::create()` to construct a ring buffer map programmatically from userspace.

– Exact API name given but limited mechanism detail aya-build-v0.2.1

THINNER COVERAGE BELOW

05 Decoupled PerfEventArray API and perf counter reads IMPROVED

55

Decouples the PerfEventArray perf buffer API and adds support for reading perf counters.

– Named API but no detail on new read interface aya-build-v0.2.1

06 Lazy BTF loading and on-demand kernel feature probing

IMPROVED

45

Custom BTF sources are now lazily loaded to reduce unnecessary I/O at startup, and kernel feature probing happens on demand rather than eagerly at load time.

– Behavioral change explained but no API surface aya-build-v0.2.1

07 Linux bindings updated to libbpf v1.7.0 IMPROVED

35

Updates Linux bindings to libbpf v1.7.0.

– Specific version but no further detail aya-build-v0.2.1

08 eBPF program loading with ksyms NEW

30

Supports loading eBPF programs with ksyms.

– Bare capability statement with no mechanism aya-build-v0.2.1

↳ BREAKING ON UPGRADE

! Map methods in aya-ebpf and userspace no longer accept types via Borrow traits — callers that relied on Borrow-based generics on map get/insert/remove methods must pass the concrete key/value types directly.

WAS THIS USEFUL?



security-profiles-operator

Feedback



v1.1.1

NOTES >

RANK

The Kubernetes Security Profiles Operator

security-profiles-operator's v1.1.1 release focuses entirely on supply-chain hardening, adding SLSA Build L3 provenance, signed/attested OCI artifacts, reproducible builds, and new seccomp base profiles for kubernix.

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

01 Supply chain hardening for release artifacts NEW 73

v1.1.1 generates SLSA Build L3 provenance for all release assets via GitHub Actions, publishes `spoc` and the Helm chart as OCI artifacts to `registry.k8s.io/security-profiles-operator` with GitHub provenance, and signs and attests platform images in the manifest list repository, including signed SLSA verification summaries from the image promoter. Per-arch container images are now reproducible, every staging artifact the promoter checks is attested and verified, and release SBOMs are attested.

– Names SLSA level, registry path, and artifacts but no verification command. v1.1.1

THINNER COVERAGE BELOW

02 Seccomp base profiles recorded on kubernix NEW 30

Records seccomp base profiles on kubernix for `amd64` and `arm64`.

– Bare statement with architectures named, no further detail. v1.1.1

WAS THIS USEFUL?



DETECTION & RESPONSE

rustinel

[Feedback](#)

v1.10.0

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

Endpoint detection for Windows, Linux, and macOS. Sigma, YARA, and IOC rules on native telemetry.

rustinel v1.10.0 and v1.9.1 add sensor event-loss alerting and stable host/agent identity across all alert types, local Sigma/YARA rule overlays and unprivileged integration access without root, plus a string of Windows, macOS, and doctor/sigma-doctor reliability fixes — several of which change default behaviour or exit codes.

←→ WHAT SHIPPED · 12 FEATURES most completely described first ? what's the number?

01 Sensor event-loss alerts for dropped telemetry NEW

90

A new `telemetry.loss_alert_interval_secs` config key (default `60`, `0` disables) emits an `agent::loss::sensor_events` alert per telemetry channel/category that dropped events in the last window, including the channel, category, and dropped count; these alerts bypass deduplication and never trigger active response. They are emitted by default (including one at shutdown), so downstream pipelines should route or filter on `event.dataset: "edr.agent"` or disable the interval.

– Names config key, alert name and filter field v1.10.0

02 Stable host and agent identity on all alerts NEW

90

A new `[agent] directory` config option persists `agent.id` as a random UUID in `<agent.directory>/agent-id`; Sigma, YARA, IOC, and correlation alerts now carry `host.id`, `host.name`, `agent.id`, `agent.type`, and `agent.version` in both NDJSON output and webhook bodies. Cloned or imaged hosts must not copy the `agent-id` file or they will share an identity.

– Config key, file path and field names all given v1.10.0

03 Local rule overlays for Sigma and YARA NEW

85

New `scanner.sigma_local_rules_paths` and `scanner.yara_local_rules_paths` config keys load local rule directories alongside a managed pack, with local rules taking precedence over pack rules of the same Sigma `id` or YARA name.

– Two exact config keys and precedence rule given v1.9.1

04 Unprivileged integration access to config and rules NEW

A new `security.integration_group` config key lets an unprivileged integration edit configuration and read logs via a Unix group or Windows service account without requiring root, and `security.integration_rules_directory` delegates control of a single rules directory to that same group or account.

– Two exact config keys and mechanism named `v1.9.1`

05 More honest 'sigma doctor' verdicts IMPROVED 83

`sigma doctor` gains `summary.can_fire_with_caveats` in its `--json` output and now exits `0` when every rule can fire (with or without caveats), reporting reliably-populated `process_creation` fields as caveats instead of marking rules degraded. Scripts relying on the old non-zero exit code for `process_creation` rulesets must be reviewed.

– Exact subcommand, field name and exit-code change given `v1.10.0`

06 Telemetry snapshot health check in 'rustinel doctor' BREAKING 76

`rustinel doctor` now reports missing, unreadable, corrupt, or stale telemetry snapshots as a failure (exit `1`) and adds a `telemetry_snapshot` entry to `--json` output; monitoring that treats any nonzero exit as a failure will begin alerting.

– Exact exit code and JSON field named `v1.9.1`

07 Richer Windows capture context IMPROVED 73

Windows PowerShell script-block alerts now include `edr.powershell.message_number` and `edr.powershell.message_total` fields so script fragments can be distinguished from whole scripts, and Windows captures now preserve the acting process ID and image for registry, file, DNS, PowerShell, and image-load events.

– Named fields and event types, no config/flag to act on `v1.9.1`

08 Overwrite protection for capture and replay output BREAKING 70

`capture` and `replay --output` now refuse to overwrite an existing capture payload, manifest, or replay report; scripts that reuse output locations must remove or rename the old files or choose a new path.

– Exact subcommand and flag named with migration note `v1.9.1`

09 Loss-rate thresholds in 'rustinel doctor' IMPROVED 65

`rustinel doctor` now evaluates drop counters as a loss rate rather than a bare count, warning from 0.01% and failing with an alert-loss result from 0.1%; pipeline queue and artifact resolver thresholds changed accordingly.

– Exact thresholds given but no exact flag/output field `v1.10.0`

10 Linux RLIMIT_NOFILE raise for per-CPU tracepoints IMPROVED

65

On Linux, the sensor now raises its soft `RLIMIT_NOFILE` to the hard limit at startup to support a per-CPU tracepoint fix; if descriptors cannot be opened, `rustinel doctor` reports the feature as degraded.

– Named limit and doctor behaviour, no runnable step `v1.10.0`

11 Windows memory scans cover cached/write-combine pages

IMPROVED

60

Windows memory scans now include pages carrying modifier bits such as `PAGE_NOCACHE` or `PAGE_WRITECOMBINE`, closing a gap in memory-based detections.

– Named bit flags but no user-facing action `v1.9.1`

THINNER COVERAGE BELOW

12 macOS network capture resilience IMPROVED

55

macOS network capture picks up new interfaces without a restart and automatically restarts failed BPF workers with backoff.

– Behaviour described, no named surface to act on `v1.9.1`

↳ BREAKING ON UPGRADE

- ! `sigma doctor` now exits `0` when every rule can fire (with or without caveats); scripts that relied on the old non-zero exit code on `process_creation` rulesets must be reviewed.
- ! Sensor event-loss alerts (`agent::loss::sensor_events`) are emitted by default (including one at shutdown); downstream pipelines should route or filter on `event.dataset:` `"edr.agent"`, or set `telemetry.loss_alert_interval_secs = 0` to disable.
- ! All alerts now include `host.*` and `agent.*` fields; `agent.id` is persisted in `<agent.directory>/agent-id` — cloned or imaged hosts must not copy this file or they will share an identity.
- ! `rustinel doctor` now evaluates drop counters as a loss rate (warn from 0.01%, alert-loss fail from 0.1%) rather than a bare count; pipeline queues and the artifact resolver thresholds have changed accordingly.
- ! On Linux, the sensor now raises its soft `RLIMIT_NOFILE` to the hard limit at startup to support the per-CPU tracepoint fix; if descriptors cannot be opened, `rustinel doctor` reports the feature as degraded.
- ! `capture` and `replay --output` now refuse to overwrite an existing capture payload, manifest, or replay report — scripts that reuse output locations must remove or rename the

old files or choose a new path.

! `rustinel doctor` now exits 1 when the service is running but the telemetry snapshot is missing or stale; monitoring that treats any nonzero exit as a failure will begin alerting.

COVERAGE

NEW detections

[Classic PowerShell start events \(ps_classic_start\)](#), [Windows Application event log channel](#), [Windows Defender Operational channel events](#), [WMI permanent subscription events via ETW](#)

NEW supported platforms

[YARA written-file scanning on macOS and Windows](#), [Linux container path resolution for process artifacts](#)

[All of rustinel's coverage changes →](#)

v1.9.0

NOTES ↗

STABLE

Adds all-matches Sigma mode (`scanner.sigma_match_mode`), per-rule compatibility diagnostics, signed update verification, and `rustinel doctor` gap warnings.

↳ GET THIS VERSION

```
$ git clone --branch v1.9.0 https://github.com/Karib0u/rustinel.git
# already have the repo? check out this version:
$ git checkout v1.9.0
```

- › Adds `scanner.sigma_match_mode = "all"` to emit every matching Sigma detection rule instead of the first.
- › Adds per-rule Sigma compatibility diagnostics to identify unsupported fields.
- › Binary updates now verify signed release checksums and rule-pack installs verify signed catalogs.
- › Adds `rustinel doctor` warnings when dropped written-file jobs create a YARA and hash IOC detection gap.

↳ BREAKING ON UPGRADE

- ! Unknown configuration sections or options, including `EDR__` environment variables, and invalid active-response settings now stop startup — remove obsolete keys before upgrading.
- ! Configuration and rule inputs writable by untrusted accounts now stop startup or leave the affected detector unloaded; check paths with `rustinel doctor`.

- ! YARA rules without recognized severity metadata now default to `high` instead of `critical`; deployments with `response.min_severity = "critical"` must add `severity = "critical"` to rules intended to trigger response, or lower the threshold.
- ! Custom rule catalogs must include a valid signature from the trusted release key; missing or invalid signatures stop replacement.
- ! The field-availability contract now uses schema version 3 with an explicit field view — update consumers that validate this schema.
- ! Log and capture output permissions are now restricted to their owner.

WAS THIS USEFUL?



Trend Vision One

[Feedback](#)[Helm Chart 3.5.5 \(2026-10-14\)](#)[BLOG](#)[RANK](#)

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

Trend Vision One's container security [Helm](#) chart now warns on unsupported chart versions during install or upgrade and bumps its [falco gRPC](#) dependency.

[WHAT SHIPPED](#) • [2 FEATURES](#) most completely described first [? what's the number?](#)

01 Unsupported chart version warning

NEW

60

The Helm chart adds a warning shown during `helm install` or `helm upgrade` when an unsupported chart version is detected, helping operators catch outdated deployments before they apply.

Upgrade the Trend Vision One container security Helm chart via Public [ECR](#) to receive unsupported-version warnings before the chart applies.

```
$ helm upgrade \
  --values overrides.yaml \
  --namespace trendmicro-system \
  trendmicro \
  oci://public.ecr.aws/trendmicro/container-security/trendmicro-
  container-security
```

– Names exact commands but not detection mechanism or thresholds

[Helm Chart 3.5.5 \(2026-10-14\)](#)

THINNER COVERAGE BELOW

02 Falco gRPC dependency upgrade

IMPROVED

50

Upgrades the falco gRPC dependency from `v1.44.0` to `v1.67.1`.

– Exact version numbers given but no behavioral impact described

[Helm Chart 3.5.5 \(2026-10-14\)](#)

WAS THIS USEFUL?



Fast and efficient osquery management

osctrl v0.6.0 introduces opt-in vulnerability monitoring and real-time SSE updates across operator workflows, alongside a generic HTTP log sink, expanded UI localization, and several configuration and performance refinements.

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

01 **SSE live updates across operator workflows** NEW 70

Adds Server-Sent Events (SSE) live-update support across all operator workflows: live query and carve updates, alert rule/channel/history changes, coalesced dashboard and node-list fleet notifications, console, file explorer, and API/TLS-to-frontend update hints. Includes a new SSE enabled parameter in the development environment configuration.

– Lists mechanism and surfaces but no exact config key name v0.6.0

02 **`read\_max` flag for osquery read buffer** NEW 65

Adds `read_max` flag to raise osquery's read buffer to 1 GB.

– Names exact flag and value but no further mechanism v0.6.0

THINNER COVERAGE BELOW

03 **Opt-in vulnerability monitoring** NEW 57

Adds opt-in vulnerability monitoring across backend and frontend, including NVD CPE matching, posture evidence, vulnerability alerts, escalation alerts, feed health tracking, hardening, and metrics.

– Names components but no config key or toggle given v0.6.0

04 **20-language UI localization with RTL support** IMPROVED 55

Expands UI localization to 20 languages, adding Farsi, Arabic, Hindi, Hebrew, Turkish, Ukrainian, and Greek, with RTL support.

– Lists concrete languages but no UI path given v0.6.0

05 **Generic HTTP log sink** NEW 45

Adds a generic HTTP log sink for forwarding osctrl logs to any HTTP endpoint.

– Describes capability but no endpoint path or config detail v0.6.0

06 Environment and service configuration controls IMPROVED

40

Makes inactive-node threshold configurable per environment, and adds per-field service-config overrides allowing service-config overrides to be pinned per field.

– Describes behavior but no flag or config key named `v0.6.0`

07 Interactive configuration prompt for `osctrl-cli` NEW

40

Adds an interactive configuration prompt for `osctrl-cli`.

– Names the CLI but no command or flow detail `v0.6.0`

08 Database performance and log attribution improvements

IMPROVED

32

Adds hot-path database indexes and attributes osquery log batches to the authenticated node.

– Backend improvement with no measurable detail given `v0.6.0`

09 `osctrl-mcp` build metadata and logging identification IMPROVED

30

Adds build metadata in `osctrl-mcp` and identifies `osctrl-mcp` in API logs.

– Thin operational change with minimal detail `v0.6.0`

WAS THIS USEFUL?



OpenCTI

[Feedback](#)

7.261016.0

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

Open Cyber Threat Intelligence Platform

OpenCTI's latest release adds entity merging from the knowledge view, defanged-indicator search, and a one-click Attack Pattern Matrix toggle, alongside conditional Fintel report sections, minor UI layout tweaks, and a new config key to cap ZIP artifact extraction size.

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

01 ZIP extraction size limit for artifact uploads NEW

90

Adds the `app:artifact_zip_max_extracted_size` / `APP__ARTIFACT_ZIP_MAX_EXTRACTED_SIZE` config key to set the maximum extracted size (in bytes) of an archive uploaded as an artifact; archives exceeding the limit are imported as-is rather than extracted.

Cap ZIP extraction at 500 MB to prevent runaway disk usage when analysts upload large archive artifacts.

```
yaml
app:
  artifact_zip_max_extracted_size: 524288000
```

– Named config key, env var, and example value given. [product docs](#)

02 Defanged indicator search in global search NEW

65

Global search filters now support `defanged` as a search term, letting analysts search using defanged indicators such as `hxxp://evil[.]com` which resolve automatically against their fanged equivalents.

Find defanged indicators (e.g., `hxxp://evil[.]com`) in global search without manual reformatting.

📌 In the OpenCTI top-bar global search, type a defanged indicator such as `hxxp://evil[.]com` — OpenCTI now resolves it against fanged equivalents automatically.

– Names the search behavior and gives a concrete example. [7.261016.0](#)

03 Expand/collapse for Attack Pattern Matrix NEW

65

The MITRE ATT&CK Attack Pattern Matrix in an Intrusion Set or Campaign's Knowledge view now supports expanding or collapsing all technique rows at once via **Expand all** / **Collapse all** controls.

Collapse or expand the full MITRE ATT&CK matrix in one click to focus on relevant techniques.

 In OpenCTI, open an Intrusion Set or Campaign > Knowledge > Attack Pattern Matrix, then click **Expand all** or **Collapse all** to toggle all technique rows at once.

– Exact UI navigation and controls provided, limited further detail. 7.261016.0

THINNER COVERAGE BELOW

04 **Conditional HTML sections in Fintel reports** IMPROVED

 45

Generated Fintel reports now automatically remove HTML sections when their template variables are empty, avoiding blank sections appearing in the output.

– Describes the behavior but no template syntax or config shown. 7.261016.0

05 **Entity merging from knowledge view** NEW

 35

Entities can now be merged directly from the knowledge view rather than through a separate workflow.

– One-line description, no mechanism or exact UI path given. 7.261016.0

06 **Description field layout on key entities** IMPROVED

 25

Improves the layout of the description field on key objects such as Intrusion Sets, Malware, and Tools.

– Bare mention of a UI tweak with no specifics. 7.261016.0

WAS THIS USEFUL?



Databricks

[Feedback](#)


TypeSafe AI Jev 1.14.0 is now available on Unity Gateway (2026-10-16)

[BLOG](#)

RANK

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

Databricks made external Apache Spark writes to Unity Catalog managed Delta tables and the SQL-target agentic code converter generally available, rounded out by GA for Unity Gateway consumer access, new Genie Agents and Genie One capabilities, and CJK full-text search support in Lakebase.

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

01 CJK full-text search in Lakebase

NEW

95

`lakebase_tokenizer` 0.2.0 adds Chinese, Japanese, and Korean full-text search via the built-in `tokenizer_chinese`, `tokenizer_japanese`, and `tokenizer_korean` configurations, plus a `tokenizer_morph` template that uses morphological analysis with bundled language dictionaries to segment CJK text, supporting compound-word modes (whole, split, or both), Japanese particle/function-word filtering and base-form conversion, and Korean particle/ending filtering. Output is compatible with GIN indexes, the `lakebase_text` index, and the `lakebase_bm25` index for BM25 ranking, so existing index infrastructure works unchanged; existing installations upgrade via `ALTER EXTENSION lakebase_tokenizer UPDATE;` without rebuilding configurations or indexes.

Upgrade an existing `lakebase_tokenizer` installation to 0.2.0 to gain CJK support without rebuilding existing configurations or indexes.

```
$ ALTER EXTENSION lakebase_tokenizer UPDATE;
```

– Names exact tokenizer configs and gives the upgrade command

Lakebase Search supports Chinese, Japan...

02 Agentic code converter GA for SQL conversion

NEW

70

The agentic code converter is now generally available for the SQL target framework, converting T-SQL, Snowflake, Redshift, Oracle, BigQuery, and Teradata dialects into Databricks SQL queries. Conversions are organized within a migration project that validates output and writes results to a workspace folder.

– Names six source dialects but gives no CLI or API syntax

The agentic code converter is generally...

03 External writes to Unity Catalog managed Delta tables **NEW**

60

External Apache Spark clients can now create and write to Unity Catalog managed Delta tables, generally available, requiring catalog commits to be enabled on the target tables.

– Names the catalog-commits requirement but no exact config steps

External writes to managed Delta tables...

04 Collaboration and context awareness in Genie One **NEW**

60

Genie One now supports collaborative document editing — users with **CAN EDIT** permission can edit shared documents — returns a user's group memberships alongside user info so workspace instructions and table descriptions can tailor responses per group, and can reuse context from past investigations to answer similar questions faster.
:

– Names the CAN EDIT permission but lacks API or config detail

Genie One enhancements (2026-10-15)

THINNER COVERAGE BELOW

05 Consumer access to Unity Gateway services **IMPROVED**

50

Consumer access to query Unity Gateway shared model and MCP services is now generally available, removing the need for an account-admin preview toggle. Account admins can set budgets and rate limits, or disable direct model access, for consumer access users to control increased model traffic to external providers.

– Describes admin controls but no concrete config surface

Consumer access to query Unity Gateway ...

06 Flash mode and upgraded models in Genie Agents **NEW**

45

Genie Agent mode adds a Flash (Beta) option that prioritizes response speed over depth for simpler, everyday tasks, alongside upgraded underlying models for better long-context handling and faster responses.

– Names the Flash beta option but no setup steps given

Genie Agents enhancements (2026-10-15)

COVERAGE

NEW integrations

[TypeSafe AI Jev 1.14.0 on Unity Gateway](#), [Power BI .tmdl file import via /importBI Genie Code skill](#), [Integrated CDC pipeline for MySQL — generally available](#), [Postgres logical replication for Lakebase](#)

NEW supported platform

[Lakehouse Real-Time SQL warehouse support for AI/BI dashboards](#)

[All of Databricks's coverage changes →](#)

TypeSafe AI Jev 1.14.0 is now available on Unity Gateway (2026-10-16)

BLOG ↗

LATEST

STABLE

Dashboard enhancements and bug fixes (2026-10-15)

BLOG ↗

STABLE

AI/BI dashboards add Genie-powered insight cards, pivot table drill-down, and theme font inheritance for rich text widgets.

- › Adds Dashboard Insights (Beta): Genie analyzes charts on a published dashboard page and surfaces trends and outliers as insight cards for viewers.
- › Adds pivot table drill-down, letting viewers traverse row and column hierarchies one level at a time.
- › Rich text widget titles and descriptions now inherit the dashboard theme's font styles.

MySQL integrated CDC pipeline is now generally available (2026-10-15)

BLOG ↗

STABLE

Replicate data in and out of Lakebase with Postgres logical replication (2026-10-15)

BLOG ↗

STABLE

Lakebase gains native Postgres logical replication, enabling live inbound migration with sync-until-cutover and outbound change publishing to external Postgres databases.

- › Supports native Postgres logical replication to stream changes into Lakebase, enabling live migration from an existing Postgres database with continuous sync until cutover.
- › Supports publishing changes from Lakebase to external Postgres databases via logical replication.

WAS THIS USEFUL?



rsigma

[Feedback](#)

v0.24.0

[NOTES ↗](#)[RANK](#)

A complete Sigma detection engineering toolkit: parser, linter, evaluator, correlation engine, conversion framework, streaming daemon, MCP and LSP servers

RSigma v0.24.0 adds parse-time semantic validation and a shared encoding-modifier expansion engine, corrects PostgreSQL and LynxDB query generation, and ships a cluster of breaking API and CLI changes that require migration.

↳ WHAT SHIPPED • 9 FEATURES most completely described first ? what's the number?

01 PostgreSQL backend query correctness fixes

BREAKING

90

The PostgreSQL backend now renders plain equality as case-insensitive `ILIKE` (as Sigma requires), renders negation as `(expr) IS NOT TRUE` to correctly handle missing fields, enforces rule ordering in `temporal_ordered` correlations via per-rule CTEs, casts JSONB `->>` text to `numeric` for numeric and comparison operators, compares booleans with `ILIKE 'true'`, tests key presence via `->` so a field whose value is JSON `null` is treated as existing, maps a plain `re` to case-sensitive `~` and `|i` to `~*`, and `|m` adds `(?w)` for line-anchor support. The `case_sensitive_re` option is removed; use the `|i` modifier on `re` values to select case-insensitive matching (`~*`) instead.

– Enumerates every fix and the removed option with replacement v0.24.0

02 Parse-time semantic validation

BREAKING

85

A new `rsigma_parser::validate` module rejects unusable logsources, invalid modifier/value combinations, malformed regexes and CIDRs, empty detections, and unresolved condition selectors before compilation. It also adds a `deprecated_detection_timeframe` lint that warns when a Sigma v1.x `timeframe:` key inside `detection:` has no effect, and a `SigmaParserError::InvalidModifiers` variant for library users catching parse errors on conflicting modifier combinations. As a result, `parse_sigma_yaml`, `parse_sigma_file`, and `parse_sigma_directory` now return parse errors for rules missing a `logsource` or containing invalid modifier/value combinations, so code that previously built such rules must be fixed.

– Names module, lint and error variant with clear before/after behaviour v0.24.0

03 Encoding modifier expansion across backends

NEW

80

A new `rsigma_ir::encoding` module is now shared by `eval` and all native backends, converting the `windash`, `wide`, `utf16le`, `utf16be`, `utf16`, `base64`, and `base64offset` modifiers into an OR of one plain match per encoded variant instead of failing with `UnsupportedModifier`. It also adds `RegexFlags::inline_prefix` to render inline regex flag groups (e.g. `(?i)`) for custom backends.

– Mechanism and full list of named modifiers given `v0.24.0`

04 Reverse conversion requires explicit logsource

BREAKING

75

`rule reverse` now exits with code `3` unless `--logsource-product`, `--logsource-category`, or `--logsource-service` is set, and `rsigma_convert` reverse conversion fails with `ConvertError::RuleConversion` when no logsource is configured.

– Exact flags, exit code and error type given `v0.24.0`

05 LynxDB backend regex and CIDR rendering fixes

IMPROVED

65

The LynxDB backend now renders regexes and CIDR as `match()` and `cidrmatch()` in-place rather than as global `| where` stages, and correctly handles `cased`, `?` wildcards, literal `*`, and values containing `/`, quotes, `>`, or mixed wildcards.

– Names concrete rendering functions and edge cases fixed `v0.24.0`

THINNER COVERAGE BELOW

06 Stricter validation errors

BREAKING

50

Lowering a detection item that applies `exists` to a keyword or to a non-boolean value now fails with `IrError::IncompatibleValue` instead of silently treating non-boolean values as `true`, and library users of `rsigma-eval` now receive `EvalError::IncompatibleValue` for CIDR values with host bits set, replacing the former `EvalError::InvalidCidr`.

– Names error type and prior silent behaviour, but library-only Names both old and new error types but no further mechanism

`v0.24.0`

07 'rule validate' accepts a single rule file

IMPROVED

45

`rule validate` now accepts a single rule file in addition to a directory.

– Exact CLI command named but minimal explanation `v0.24.0`

08 HIR schema version bump

35

BREAKING

`HIR_SCHEMA_VERSION` is now 3; consumers that check this constant must be updated.

– Bare version bump with no further detail `v0.24.0`

09 Processing-pipeline and correlation API changes

BREAKING

15

Processing-pipeline and correlation APIs have changed; library users must update call sites per the migration notes.

– Vague note with no named surfaces or specifics `v0.24.0`

↳ **BREAKING ON UPGRADE**

- ! `parse_sigma_yaml`, `parse_sigma_file`, and `parse_sigma_directory` now return parse errors for rules missing a `logsource` or containing invalid modifier/value combinations; code that previously built such rules must fix them.
- ! `HIR_SCHEMA_VERSION` is now 3; consumers that check this constant must be updated.
- ! Processing-pipeline and correlation APIs have changed; library users must update call sites per the migration notes.
- ! The PostgreSQL `case_sensitive_re` option is removed; use the `|i` modifier on `re` values to select case-insensitive matching (`~*`).
- ! `rule reverse` now exits with code 3 when none of `--logsource-product`, `--logsource-category`, or `--logsource-service` is set.
- ! Lowering a detection item that applies `exists` to a keyword or to a non-boolean value now fails with `IrError::IncompatibleValue` instead of silently treating non-boolean values as `true`.
- ! `rsigma_convert` reverse conversion fails with `ConvertError::RuleConversion` when no logsource is configured.
- ! Library users of `rsigma-eval` receive `EvalError::IncompatibleValue` for CIDR values with host bits set, replacing the former `EvalError::InvalidCidr`.

WAS THIS USEFUL?



EXPOSURE & VULNERABILITY MANAGEMENT

syzkaller

Feedback



changes since 2026-10-14

CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

Patching flow gains a `patch-planner` agent and `grepper` / `gitlog` tools; agents rebalanced to `CoreModel`; Gemini backend now uses `gemini-4-argon`.

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-14
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-14
```

- › Adds `patch-planner` agent to the `pkg/aflow/flow/patching` pipeline, backed by the deep-reasoning model, to plan patches before code edits are generated.
- › Adds `grepper.Tool` and `gitlog.Tools` to the code-access tool set available to patching-flow agents (alongside the existing `codesearcher.Tools`), exposed as the `codeAccessTools` bundle in `pkg/aflow/flow/patching`.
- › Switches `history-explorer` agent from `DeepReasoningModel` to `CoreModel` in the patching flow, reducing deep-reasoning quota consumption.
- › Switches the C-repro loop and initial-researcher agents to `CoreModel` in `pkg/aflow/flow/reproc`.
- › Updates the Gemini backend to use `gemini-4-argon` as the deep-reasoning model in `pkg/aflow/backend/gemini`.

COVERAGE

NEW detections

[UFS BSG request syscall descriptions, GPIO character device syscall descriptions](#)

UPDATED detection [Binder transaction and read buffer descriptions](#)

[All of syzkaller's coverage changes →](#)

WAS THIS USEFUL?



◆ Identity & Access Management

Ory

[Feedback](#)[changelog · Oct 17, 2026](#)[CHANGELOG ↗](#)

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

Ory v26.4.x adds TraceID correlation in live event streams and finer latency histogram buckets in Ory Hydra.

- › Adds `TraceID` key in `eventAttributes` for events delivered through a live event stream, enabling correlation of all events from a single request
- › Refines the `http_requests_duration_seconds` and `http_response_time_seconds` request latency histogram buckets in Ory Hydra for finer granularity

COVERAGE

NEW

supported platform

FIPS 140-3 container images and Helm chart support for Kratos, Hydra, and Polis

All of Ory's coverage changes →

WAS THIS USEFUL?



Passbolt

[Feedback](#)[snapshot-20261017](#)[NOTES ↗](#)

Manage and share passwords securely with Passbolt. Open source, audited and built for teams that need collaboration, compliance and control.

Passbolt gains Offline Mode (beta), secret history, zero-knowledge encrypted metadata, standalone notes, PIN code resource type, custom fields, RBAC dynamic roles, and in-app edition switching.

- › Introduces Offline Mode (beta) giving users secure read-only access and autofill for selected credentials when the Passbolt server cannot be reached.
- › Introduces secret history, letting users view prior versions of stored secrets.
- › Enables administrators to run encrypted metadata in zero-knowledge mode.
- › Promotes encrypted metadata and new resource types to stable.
- › Introduces standalone notes as a distinct resource type.

[+10 more](#)

COVERAGE

NEW integrations [PingOne SSO](#), [SMTP OAuth support for Microsoft Exchange Online](#)

UPDATED integration [SCIM provisioning promoted out of beta](#)

NEW supported platforms

[Safari browser extension generally available](#), [PHP 8.5 support](#), [Slovenian and Ukrainian language support](#)

UPDATED supported platform [Microsoft Edge extension migrated to Manifest V3](#)

[All of Passbolt's coverage changes →](#)

WAS THIS USEFUL?



RANK

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

ConductorOne's biggest push this window was rolling the C1AI assistant into findings remediation, Slack, and access-review workflow steps, alongside new local approval policies, shadow apps surfacing as findings, and a batch API for campaign templates; two breaking changes affect connector-synced resource deletion and one-click request links.

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

01 Batch API and custom selectors for campaign templates NEW

85

Adds **Add template entitlements** and **Remove template entitlements** API calls, each accepting up to 200 entitlements per call, to build or trim campaign template selections in batches without overwriting earlier choices. Also introduces custom selectors for access review campaign scoping — filtering by app, resource type, risk level, and compliance framework — so new matching entitlements are pulled in automatically when a campaign is prepared.

— Names exact API calls, their limits, and scoping filter criteria

October 16, 2026 (2026-10-17)

02 Connector-synced resource deletion now errors; adds **isManuallyManaged** field

85

BREAKING

DELETE

`/api/v1/apps/{app_id}/resource_types/{resource_type_id}/resources/{id}` now returns an error when called on connector-synced resources instead of silently deleting them — previously deleted resources simply reappeared on the next sync. Callers that relied on deleting synced resources must update their logic. App resources returned by the API now also include a read-only **isManuallyManaged** field.

Verify that deleting a connector-synced resource now surfaces an error instead of silently removing it.

```
$ curl -X DELETE
'https://<tenant>.conductor.one/api/v1/apps/<app_id>/resource_types/<resource_type_id>/resources/<id>' \
-H 'Authorization: Bearer <token>'
```

– Exact endpoint, before/after behaviour, and a runnable curl example

October 16, 2026 (2026-10-17)

03 Local policies scoped to a single app or entitlement NEW

80

Adds local policies — policies bound to a single app or entitlement setting (request policy, revoke policy, or preferred review policy) that never appear in the global policy library. They are created via [Create new local policy](#) from the policy menu on any app or entitlement setting; once created their scope cannot be changed, they cannot be reused, and only Super Administrators can create or manage them.

:

Apply a one-off approval workflow to a single sensitive entitlement without polluting the global policy library.

1. Open the entitlement that needs a unique policy.
2. Go to the [Request policy](#) setting.
3. Open the policy menu and click [Create new local policy](#).
4. Configure the policy steps and save — it is bound to that entitlement immediately.
5. To edit later, return to the same [Request policy](#) setting and click [Edit local policy](#).

– Names exact menu action, scope rules, and permission restriction

October 16, 2026 (2026-10-17)

04 New findings capabilities NEW

80

Adds new findings capabilities: a [Remediate with AI](#) action, available from a finding's row menu and the finding details drawer, that opens C1AI with the finding's type and target pre-filled and follows built-in playbooks for unowned service accounts, non-human identities, shadow apps, and connector sync failures; a per-application Findings tab listing open and in-progress findings with target and assignee, with filters for severity, type, and assignee; and shadow apps are now surfaced as findings, enabling assign, snooze, suppress, accept-risk, reopen, and [Remediate with AI](#) actions directly on discovered shadow apps.

:

Get AI-guided fix steps for a specific finding without leaving the ConductorOne console.

1. In the Findings list, open a finding's row menu and select [Remediate with AI](#) —

C1AI opens with the finding's details already filled in.

Start AI-guided remediation for a finding such as an unowned service account without manually looking up remediation steps.

📌 In the Findings list, open a finding (e.g. an unowned service account), then click **Remediate with AI**. C1AI opens with the finding type and target pre-filled and walks through the built-in remediation playbook.

Access the C1AI assistant directly within ConductorOne for guided setup and operational help.

📌 In ConductorOne, locate and open the C1AI assistant from the product interface.

Review all open findings scoped to a single application, filtered by severity or assignee, during an access review.

📌 Navigate to the application's details page, select the Findings tab, then use the severity, type, or assignee filters to narrow the list.

– Names the action, trigger points, and specific remediation playbooks Clear navigation path and named filters, but limited mechanism detail Names the available actions but no steps to trigger them

October 16, 2026 (2026-10-17)

05 C1AI as automated workflow step in review policies **NEW**

60

C1AI can now be added as a workflow step in access request or review policies — selecting **C1AI** as the assignee lets the agent automatically evaluate each request that reaches that step, optionally following custom instructions entered for that step.

Enable C1AI as a workflow step in an access review policy so the agent automatically evaluates each request that reaches it.

📌 In the C1 console, open a request or review policy, add a new workflow step, and select **C1AI** as the assignee. Optionally enter custom instructions the agent should follow when evaluating requests at that step.

– Explains mechanism and gives exact UI steps to configure it [product docs](#)

THINNER COVERAGE BELOW

06 C1AI assistant natively in Slack **NEW**

55

The C1AI assistant is now available natively in the C1 app for Slack, with buttons, forms, and pickers so users can complete access tasks without leaving Slack; users can ask the

assistant what actions are available and request them directly without navigating to the Actions catalog.

When you need to request access without leaving Slack, ask the C1AI assistant directly instead of opening the Actions catalog.

 In the C1 app for Slack, open a conversation with the C1AI assistant and ask: 'What actions are available to me?' then follow the prompts to submit a request.

– Gives an exact Slack prompt but little on the underlying mechanism [product docs](#)

07 Access request links no longer auto-submit

BREAKING

55

Access request links — including one-click links from Copilot or Jira — no longer auto-submit requests; they now only prefill the form, requiring an explicit user submit. Jira templates, runbooks, or automations that expected auto-submit behaviour must be updated.

– Clear before/after and migration note, but no named config surface

October 16, 2026 (2026-10-17)

08 Unlinked accounts in high-risk role grant queries and exports

IMPROVED

45

High-risk role grant queries now include accounts not linked to an identity user — for both permanent and time-limited grants — with blank User and email cells, and these unlinked accounts are also included in CSV exports of high-risk role grants with blank User and email cells.

– Describes behaviour precisely but gives no steps to view it [product docs](#)

09 Validate-scope report adds review policy detail

IMPROVED

40

The validate-scope report's **User grants** sheet now includes the review policy, first approval step type, and expected reviewers for each review task.

– Names the sheet and new columns but no retrieval steps [October 16, 2026 \(2026-10-17\)](#)

10 App catalog search matches entitlement descriptions

IMPROVED

35

App catalog search now matches entitlement descriptions, letting users find access by describing what it does rather than knowing the exact name.

– States the change but no example query or UI path [product docs](#)

11 Character limits on request form text fields

NEW

35

Adds optional character-limit settings to Text and Text area fields in request forms, preventing submission of over-length answers.

– Names the field types but no config path given October 16, 2026 (2026-10-17)

12 Custom form responses shown in Slack and Teams notifications

IMPROVED

30

Approvers now see custom form responses, including clickable links such as Jira tickets, directly in Slack and Microsoft Teams notifications.

– States the change with no steps or configuration detail October 16, 2026 (2026-10-17)

13 Entitlement details drawer in access review table

IMPROVED

25

Reviewers can open entitlement or resource details in a drawer from within an access review table without leaving the review.

– Brief description of a UI convenience with no further mechanism

October 16, 2026 (2026-10-17)

↳ ALSO FROM THESE RELEASES

.

↳ BREAKING ON UPGRADE

- ! `DELETE /api/v1/apps/{app_id}/resource_types/{resource_type_id}/resources/{id}` now returns an error when called on connector-synced resources instead of deleting them – callers that relied on deleting synced resources will receive errors and must update their logic.
- ! Access request links (including one-click links from Copilot or Jira) no longer auto-submit requests – they now only prefill the form. Jira templates, runbooks, or automations that expected a link to create a request automatically must be updated to require an explicit user submit.
- ! `DELETE /api/v1/apps/{app_id}/resource_types/{resource_type_id}/resources/{id}` now returns an error when called on connector-synced resources; previously it deleted them and they reappeared on the next sync.
- ! One-click access request links (such as those from Copilot or Jira) no longer submit requests automatically – they now only prefill the form, requiring an explicit user submit. Jira templates, runbooks, or automations that expected auto-submit behaviour must be updated.

WAS THIS USEFUL?



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

Auth0 rolled out Early Access enhancements to Universal Login, adding tenant-level phone country controls with a redesigned phone entry UI and type-specific identifier presentation options.

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

01 Phone country controls and entry UI in Universal Login NEW 75

Universal Login now supports tenant-level include/exclude lists for supported phone countries, configurable via the PATCH /api/v2/tenants/settings endpoint. It also ships an updated phone entry UI with country selection, regional formatting, a numeric keyboard, and phone number suggestions on supported devices.

- Names endpoint and lists concrete UI mechanics, but no limits given

Universal Login identifier and phone co...

02 Identifier presentation controls for Universal Login NEW 65

Adds identifier presentation controls, letting tenants choose between type-specific Continue with buttons or a unified field, configurable via the PATCH /api/v2/branding/theme endpoint.

- Names exact endpoint and option but no further behavior detail

Universal Login identifier and phone co...

ALSO FROM THESE RELEASES

- .
- .
- .
- .

WAS THIS USEFUL?

Omada Identity

Feedback



14th October 2026 (2026-10-15)

BLOG ↗

RANK

Omada Identity is an identity governance and administration platform that manages user access, entitlements, compliance certifications, and access requests across enterprise systems.

Omada Identity shipped a small monitoring improvement to its Operations dashboard, adding instance identification for dual-CAG environments.

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

01 Instance ID shown in System Heartbeats widget IMPROVED

53

The System Heartbeats widget in the Operations dashboard now displays the instance ID next to each worker service, letting operators distinguish between customer-hosted and Omada-hosted Cloud Application Gateway instances.

– Names the exact widget and dashboard; no further mechanism given.

14th October 2026 (2026-10-15)

WAS THIS USEFUL?



Cyera

Feedback



announced · Oct 16, 2026

PRESS ↗

RANK

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

Cyera launched Agent Guardian in the Claude Marketplace, giving security teams real-time visibility and policy enforcement over Claude agents.

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

01 Agent Guardian in Claude Marketplace NEW 65

Cyera's Agent Guardian is now available in the Claude Marketplace, providing continuous real-time visibility into every Claude agent in use — including sanctioned and shadow agents — covering tools, dependencies, and data connections. It can intercept out-of-policy commands, unapproved database changes, and unauthorized data transfers before they execute, with guardrails anchored in data classification to keep autonomous Claude agents within approved business boundaries. Listing in the marketplace also lets Anthropic customers apply committed Anthropic spend toward Cyera licensing.

– Explains mechanism and scope but no UI path or config detail launch-20261016-52f56078

WAS THIS USEFUL?



Twingate

Feedback



Privileged Access for Web Apps: HTTPS

NOTES ↗

RANK

Twingate is a identity-based access for users, services, and AI agents that deploys in minutes, scales to every resource, and finally lets you retire your VPN.

Twingate extended Privileged Access for Web Apps to support HTTPS, delivering end-to-end encrypted access to internal web apps without certificate warnings or extra sign-ins.

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

01 HTTPS support for Privileged Access for Web Apps NEW

63

Privileged Access for Web Apps now covers HTTPS-served internal web apps across every leg of the user's path, including the final hop to the origin. This eliminates browser **Not Secure** warnings and secondary sign-in prompts, supports both public and private CAs, and uses group-driven access control with no per-app certificate management required.

– Explains mechanism and scope but no config keys or commands given.

Privileged Access for Web Apps: HTTPS

WAS THIS USEFUL?



◆ Non-Human & Agent Identity

warden

[Feedback](#)[v0.21.0](#)[NOTES ↗](#)[RANK](#)

The secure gateway connecting AI agents to enterprise systems.

Warden v0.21.0 ships a one-command dev playground for trying out governed AI-agent workflows, expands keyless credential federation to Anthropic, OpenAI, and Cloudflare, introduces an [MCP Skills](#) extension for identity-bound skill discovery, and switches the default assertion profile to an [RFC 8693](#) delegation token — a breaking change for any claim-reading verifier.

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

01 One-command dev playground and scenario tour NEW 95

`warden server -dev-playground` spins up a self-contained Warden instance with its own [identity provider](#), a mock bank exposing MCP and REST interfaces, and GitHub's MCP server, with no upstream key or Compose stack required. New `warden dev jwt`, `warden dev scenarios`, and `warden dev audit` subcommands drive a nine-scenario tour across seven agent clients (Claude Code, Codex, Cursor, Gemini CLI, opencode, VS Code, or any MCP client), runnable interactively via `warden dev scenarios -client`.

Instantly start a governed AI-agent environment with no pre-existing identity provider or secrets — ideal for a first hands-on demo.

```
$ warden server -dev-playground
```

Walk through all nine governance scenarios interactively using your own MCP-capable agent to see policy enforcement in action.

```
$ warden dev scenarios -client
```

— Exact commands, subcommands, scope, and supported agents given. [v0.21.0](#)

02 Keyless credential federation expansion NEW 90

Adds `warden cred source|spec keyless-plan` to describe the keyless replacement for a given credential source or spec, a `keyless_enforcement_level` config key to warn or refuse writes that would leave a secret stored in Warden, and a `stored_secrets` field on every source and spec. Adds keyless [workload identity](#) federation for Anthropic and OpenAI providers, a keyless Cloudflare credential source, and Azure Key Vault as a chaining producer for keyless credentials. Azure

`key_vault_secret` specs now fail every mint and must be migrated to `secret_read`.

Check which credentials on a source would need to move to keyless before enforcing `keyless_enforcement_level`.

```
$ warden cred source keyless-plan
```

– Names config key, commands, providers, and migration path. v0.21.0

03 Assertion profile control and RFC 8693 default

BREAKING

85

Adds an `assertion_profile` config key accepting `default`, `minimal`, or `aws`, plus a per-spec `assertion_ttl`, to control the shape and lifetime of federated assertions. The `default` profile is now an RFC 8693 delegation token placing the user at the top level and the agent under `act`, removing the `warden_sub`, `warden_auth_mount`, `warden_user`, and `warden_namespace` claims — any verifier (Vault `JWT` role, GCP attribute condition, `CEL` rule) reading those claims must be rewritten. New AWS federated specs default to the `aws` profile (requiring `sts:TagSession`); new Azure federated specs default to `minimal`.

– Names config keys, removed claims, and default-profile impact. v0.21.0

04 MCP Skills extension for identity-bound skill discovery

85

BREAKING

Adds the MCP Skills extension for skill discovery bound to agent identity, including `read_skill(uri)` to retrieve a skill's markdown via a `skill://` URI in `structured output`, replacing the removed `get_skill`. `list_roles` now returns structured `skill` and `provider_path` fields with the resolved provider URL and `skill://` URI alongside each role; roles without a backfilled `provider_path` list with no URL. Stored skill names containing underscores are renamed at unseal — e.g. `mcp_aws` becomes `mcp-aws`, `ansible_tower` becomes `ansible-tower`.

– Names functions, fields, URI scheme, and rename examples. v0.21.0

05 JSON-RPC error format for MCP refusals

BREAKING

75

MCP refusals are now returned as JSON-RPC errors with code `-32090`; the former `error` / `error_description` fields move under `error.data`.

– Names exact error code and field relocation for migration. v0.21.0

06 GitHub App token scoping by repository and permissions

BREAKING

GitHub App token scoping now uses `repositories` and `permissions` fields; the former singular `repository` field is refused, and a stored `permissions` value is now enforced.

– Names exact fields and the breaking migration requirement. `v0.21.0`

THINNER COVERAGE BELOW

07 Merge semantics for AWS and Alibaba Cloud config writes**IMPROVED**

45

AWS and Alibaba Cloud config writes now merge — a key omitted from a write retains its existing value rather than being cleared.

– Describes behavior change but no config key names given. `v0.21.0`

08 Public token_exchange client support**NEW**

30

Adds support for public `token_exchange` clients.

– Bare mention, no mechanism or scope described. `v0.21.0`

↳ BREAKING ON UPGRADE

- ! The `default` assertion profile is now an RFC 8693 delegation token: with a user disclosed, the top level becomes the user and the agent moves into `act`. The claims `warden_sub`, `warden_auth_mount`, `warden_user`, and the agent-level `warden_namespace` are removed — any verifier (Vault JWT role, GCP attribute condition, CEL rule) reading those claims must be rewritten.
- ! `get_skill` is removed; callers must migrate to `read_skill(uri)`. The (skill: ..., url: ...) description convention is replaced by the role fields `skill` and `provider_path` — roles without `provider_path` backfilled will be listed with no URL.
- ! Skill reads on the discovery server are now identity-bound.
- ! Stored skill names containing underscores are renamed at unseal — e.g. `mcp_aws` becomes `mcp-aws`, `ansible_tower` becomes `ansible-tower`.
- ! MCP refusals are now JSON-RPC errors (`-32090`); the former `error` / `error_description` fields move under `error.data`.
- ! Azure `key_vault_secret` specs now fail every mint — they must be moved to `secret_read`.
- ! GitHub's singular `repository` field is refused; specs must use `repositories`. A stored `permissions` value is now enforced.
- ! New AWS federated specs default to the `aws` profile (requires `sts:TagSession`); new Azure federated specs default to `minimal`.
- ! AWS and Alibaba Cloud config writes now merge — a key omitted from a write retains its existing value rather than being cleared.

WAS THIS USEFUL?



GOVERNANCE, RISK & COMPLIANCE

CISO Assistant

Feedback



v4.1.2

NOTES ↗

RANK

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 200+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

CISO Assistant v4.1.2 adds GDPR/DPIA screening and a processor register role, new MCP tools for incident and evidence management, a unified outbound mail configuration, a beta OAuth log-forwarding sink for Microsoft Sentinel, and new Hungarian and Costa Rica compliance frameworks.

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

01 Unified outbound mail routing via MAILERS NEW 80

Adds a **MAILERS** service config to route all outbound mail (reminders, notifications) through a single configured mail service, replacing the previous per-feature mail settings.

Route all outbound CISO Assistant mail (reminders, notifications) through a single configured mailer service.

```
text
MAILERS=smtp://user:password@mail.example.com:587
```

– Named config key with a runnable example value. v4.1.2

02 MCP incident and evidence management tools NEW 75

Adds **create_evidence**, **add_incident** and **add_timeline_entry** MCP tools for uploading local files as evidence and managing incidents and timeline entries from the MCP layer. Also adds MCP support for linking task templates to incidents via **create_incident** and **update_incident** calls.

– Names all MCP tool calls but no sample invocation shown. v4.1.2

THINNER COVERAGE BELOW

03 OAuth sink for log forwarding to Microsoft Sentinel NEW 45

Adds an OAuth sink (beta) for log forwarding to support Microsoft Sentinel integration.

– Names integration target but no setup steps given. v4.1.2

- 04

Unique reference ID for folders

NEW

45

Adds a `ref_id` field, unique across all folders, to the folders model.

– Names field and uniqueness scope but no usage detail.

v4.1.2
- 05

Privacy register and DPIA enhancements

NEW

40

Supports DPIA screening and DPIA forms for GDPR compliance workflows, and adds a processor register role on the Processing register.

– Describes purpose but not form fields or role permissions.

v4.1.2
- 06

Applied control linked field enhancements

NEW

40

Enables linking a metric directly to an applied control, and adds evidences to applied control linked fields.

– States two linking additions without workflow detail.

v4.1.2
- 07

Extended result field in data wizard

NEW

40

Adds an `extended_result` field to requirement assessments in the data wizard.

– Names field and location but not its purpose or values.

v4.1.2
- 08

Hungarian and Costa Rica compliance frameworks

NEW

25

Adds new Hungarian and Costa Rica frameworks, extending the platform's library of 200+ global compliance frameworks.

– Only named in the release summary, no further detail.

v4.1.2

WAS THIS USEFUL?



AGENT

Runlayer

Feedback



1.28.1446 - Supported clients and card detection (2026-10-16)

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 1.28.1446 release gives admins full visibility into which AI Watch clients can be enforced before writing policy, and lets connector requests be approved against access a requester already has instead of creating duplicate grants.

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

01 Supported clients visibility for AI Watch policies

IMPROVED

69

The **View supported clients** dialog under **Manage → Endpoints** now lists every client AI Watch hooks can protect and enforce, including each client's minimum required version and whether it blocks before the tool runs only. The Enforce and Protect built-in policy pickers now surface this same full AI Watch supported client list when authoring policies.

Check which AI Watch clients and minimum versions are enforceable before authoring a new Enforce or Protect policy.

📌 In the Runlayer console, go to **Manage → Endpoints → View supported clients** to see every client AI Watch can protect, its minimum required version, and whether enforcement blocks before the tool runs.

– Names exact UI path and fields, but no numbers or API surface.

1.28.1446 - Supported clients and card ...

02 Connector request approval via existing shared connectors

NEW

60

When approving a new connector request, admins can now select a connector they already built and shared; if the requester already has access to it, it shows as **Has access** and approving the request grants access through that existing connector instead of creating another policy.

Approve an incoming connector request without granting duplicate access when the requester already has access to one of your shared connectors.

 In the Runlayer console, go to Inbox, open the connector request, select an existing connector the requester already has access to (shown as **Has access**), and approve — no additional policy is created.

– Clear workflow and UI path, no config keys or endpoints named.

1.28.1446 - Supported clients and card ...

WAS THIS USEFUL?



DEPLOY

microsandbox

Feedback



v0.7.8

NOTES ↗

RANK

easy, fast, programmable and local-first microVM runtime and library

Across v0.7.7 and v0.7.8, microsandbox adds runtime-owned detached jobs, an HTTP CONNECT outbound proxy, DAX filesystem support, and a batch of refinements to snapshots, networking and the CLI including a `--mount-disk` restore flag, published port ranges, and a `--tree` command display control.

↳ WHAT SHIPPED · 11 FEATURES most completely described first ? what's the number?

01 `--mount-disk` flag on `msb restore` NEW 72

The `msb restore` command gains a `--mount-disk` flag to attach a disk image when restoring a sandbox from a snapshot archive.

– Exact flag and subcommand named but no example run v0.7.8

02 `--tree` command display control in CLI NEW 65

The CLI adds `--tree` command tree depth and display controls, letting users inspect the full command hierarchy and available subcommands, e.g. `msb --tree`.

Inspect the full command hierarchy and available subcommands after the new tree display controls were added.

```
$ msb --tree
```

– Exact flag with a runnable example command v0.7.7

THINNER COVERAGE BELOW

03 Port ranges in published ports IMPROVED 55

`msb` published ports now support port ranges, so a span of guest ports can be exposed on the host in a single declaration.

– Named config surface but no syntax example v0.7.8

04 Header-scoped secret substitution in network config IMPROVED 50

Secret substitution in network configuration is now scoped to specific header fields, preventing secrets from being injected into unintended request headers.

– Mechanism explained but no config key named v0.7.8

05 TLS certificate trust helper for applications NEW

45

An opt-in application certificate trust helper allows applications inside a microVM to trust certificates generated by the TLS interception layer.

– No config key named for enabling this helper v0.7.8

06 Runtime-owned detached jobs NEW

40

The execution runtime can now own detached jobs, managing long-running jobs independently of the calling client.

– Described only in prose, no API or flag named v0.7.8

07 Filesystem state limits for snapshots NEW

40

Snapshots can now have filesystem state limits configured, capping how much filesystem state a sandbox snapshot may consume.

– No specific config key or flag named v0.7.8

08 HTTP CONNECT outbound proxy for sandbox egress NEW

40

Sandbox network egress now supports an HTTP CONNECT outbound proxy.

– Thin description, no configuration syntax given v0.7.7

09 DAX filesystem support and optional init binary NEW

40

Adds DAX (Direct Access) filesystem support and an optional init binary.

– Named feature but no flag or config to enable it v0.7.7

10 Guest clock sync toggle NEW

25

An option now exists to disable guest clock synchronization at runtime.

– No flag or config key named v0.7.7

11 Cache cleanup and usage reports in storage management

IMPROVED

20

Storage management gains cache cleanup, usage reports, and general CLI polish.

– Vague grouping with no named commands or flags v0.7.7

WAS THIS USEFUL?



EVALUATE

◆ AI Observability & Evals

ongrid

Feedback



v0.18.2

NOTES ↗

RANK

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

ongrid's latest releases add OpenTelemetry-based observability for agents, a built-in GPU device role for fleet management, and a unified APM pipeline for Kubernetes pod metrics, alongside UI polish to the web chat sidebar.

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

01 Unified APM pipeline for Kubernetes pod metrics

IMPROVED

45

Scopes Kubernetes pod metrics to capture selections in the APM view and unifies SDK ingestion and automatic capture under a single APM pipeline.

– Explains what changed but no config keys or limits given v0.18.0

02 OpenTelemetry observability for agents

NEW

40

ongrid adds OpenTelemetry observability for the ongrid agent, allowing agent behavior to be traced via standard OTel tooling.

– Names OpenTelemetry but no spans, exporters or config detailed v0.18.2

03 Built-in GPU device role for fleet management

NEW

40

Adds a built-in GPU device role for classifying and managing GPU-equipped devices in the fleet.

– Describes purpose but not configuration or classification criteria v0.18.0

04 Collapsible chat sidebar in web UI

NEW

30

Adds a dedicated collapsible chat sidebar to the ongrid web UI.

– Clear UI addition but no further mechanism described v0.18.2

WAS THIS USEFUL?



GOVERN

◆ AI Security Posture & Discovery

Zenity

Feedback



changelog · Oct 17, 2026

CHANGELOG ↗

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

Zenity adds Universal Runtime Control for Claude Enterprise via Anthropic's Inference Hooks to validate tool calls before execution.

- › Supports Anthropic's Inference Hooks to validate Claude Enterprise tool calls at runtime before they execute.

WAS THIS USEFUL?



garak

Feedback



v0.17.1

NOTES

RANK

Garak scans language models and applications for vulnerabilities including prompt injection and data leakage.

Garak v0.17.1 adds a hit@k attack success rate metric to its reporting output and gains support for Python 3.14.

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

01 Hit@k attack success rate metric

NEW

73

Adds a `hit@k` attack success rate metric, exposed as terminal output and per probe/detector pair in the final digest entry of `report.jsonl`.

– Names exact metric and report field, but no scope/threshold detail v0.17.1

THINNER COVERAGE BELOW

02 Python 3.14 support

IMPROVED

20

garak gains Python 3.14 support.

– Bare compatibility note with no further detail v0.17.1

WAS THIS USEFUL?



NVIDIA SkillSpector

Feedback



v2.13.0

NOTES ↗

RANK

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, security risks, prompt injection, data exfiltration, and supply-chain risks in Claude Code, Codex, and MCP skills before you install them.

SkillSpector v2.13.0 locks down MCP HTTP listeners to loopback and retires the Codex CLI analysis provider, while broadening malicious-pattern detection across bundled settings, shell payloads, Python data flows, binary files, and Claude plugin manifests, alongside baseline-integrity and batch-scan fixes.

↳ WHAT SHIPPED • 12 FEATURES most completely described first ? what's the number?

01 Baseline versioning and atomic generation

BREAKING

70

Exact-fingerprint baselines are now version-bound and must be regenerated after upgrading to v2.13.0. Baseline generation also uses atomic replacement; updates that cannot preserve ownership or group safely are refused, requiring a new baseline file to be written instead.

– Gives clear before/after and the required regeneration action v2.13.0

02 MCP HTTP listeners restricted to loopback

BREAKING

60

MCP HTTP listeners are now restricted to loopback connections; remote MCP deployments that previously accepted direct connections now require an authenticating proxy.

– Names the change and required mitigation, no exact config shown v2.13.0

THINNER COVERAGE BELOW

03 Detection for malicious bundled settings

NEW

58

Adds detection for bundled settings that use the short `-r` flag to load and run extra code, and flags bundled settings that run commands, redirect AI traffic, or automatically approve connected tool servers.

– Names concrete malicious patterns including the ``-r`` flag v2.13.0

04 NVIDIA Build default model updated to GLM-5.3

BREAKING

55

NVIDIA Build now defaults to GLM-5.3 after retirement of the previous default (GLM-5.2); explicit model overrides remain supported.

– Names both model versions and notes override path v2.13.0

05 Archive-member path collision handling

BREAKING

55

If a real directory ending in **!** collides with an archive-member path, that directory must now be renamed to allow both files to be fully inspected.

– Gives a concrete scenario and the exact fix needed v2.13.0

06 Improved Python secret-flow detection

IMPROVED

50

Detects sensitive data passed to **urllib** opener requests, including inline opener calls, and finds Python flows from secret reads to risky uses even when assignments and uses appear at different nesting depths.

– Names urllib opener mechanism and nesting-depth tracking v2.13.0

07 Broader executable-file detection

IMPROVED

50

Improves executable-file detection by checking complete TypeScript declaration files before treating them as non-executable, so partial or disguised runtime code is not skipped, and by recognizing more executable Python files, including **.pyw** files and extensionless launchers.

– Names TypeScript declaration files and `.pyw` extension specifically v2.13.0

08 Codex CLI analysis provider disabled

BREAKING

45

The Codex CLI analysis provider is disabled in this release; any workflow relying on it must switch to another supported provider.

– States the removal and migration need but no replacement list v2.13.0

09 Hidden instructions behind binary file prefixes

NEW

45

Inspects readable instructions hidden behind binary file prefixes while keeping ordinary PDF/GIF assets and XMP metadata from producing false alarms.

– Describes mechanism and false-positive guard, detection-only v2.13.0

10 OMS signing bundles treated as untrusted input

NEW

45

Treats OMS signing bundles as untrusted input, inspecting decoded metadata and preserving raw evidence and incomplete coverage when parsing fails.

– Names OMS bundles and failure-handling behaviour, no user action v2.13.0

11 Expanded shell execution and remote shell detection

IMPROVED

40

Detects encoded shell payloads that are decoded and executed, and adds lower-confidence warnings for multiline scripts that appear to open a remote shell, while preserving existing one-line detections.

– Explains new detection scope but it's scanner-side, not user-actionable v2.13.0

12 Claude plugin manifest hook inspection NEW

40

Inspects hooks declared in Claude plugin manifests and reports incomplete coverage when those declarations cannot be fully checked.

– Names the manifest surface but little mechanism beyond that v2.13.0

↳ BREAKING ON UPGRADE

- ! The Codex CLI analysis provider is disabled; any workflow relying on it must switch to another supported provider.
- ! MCP HTTP listeners are restricted to loopback; remote MCP deployments that previously accepted direct connections now require an authenticating proxy.
- ! Exact-fingerprint baselines are version-bound and must be regenerated after upgrading to v2.13.0.
- ! NVIDIA Build defaults to GLM-5.3 after retirement of the previous default (GLM-5.2); explicit model overrides remain supported.
- ! Baseline generation uses atomic replacement; updates that cannot preserve ownership or group safely are refused — write a new baseline file instead.
- ! If a real directory ending in ! collides with an archive-member path, that directory must be renamed to allow both files to be fully inspected.

WAS THIS USEFUL?



Independent Auditing of AI Agents. Run by human or the agent itself, to answer the most crucial question in the AI Agent Economy.

iFixAi's v4.1.0 and v4.2.0 releases add two new audit commands (`trajectory` and `policy-lab`), route inspections through Vercel and Cloudflare AI Gateways, and cut default run time ~6x by switching `ifixai run` to a 10-inspection suite — alongside a string of breaking changes to exit codes, scoring, and CLI validation.

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

01 Default run suite reduced from 60 to 10 inspections

BREAKING

95

`ifixai run` now defaults to the 10-inspection `essential` suite (`B01` , `B08` , `P01` , `B12` , `B09` , `B27` , `B06` , `B16` , `B18` , `B20`) instead of all 60, cutting agent and judge calls by ~6x. Use `--full-run` or `--suite` `smoke|strategic|essential|core|extended|all` to select a different suite; pipelines relying on full coverage must add `--full-run` .

– Names every inspection ID and the full suite flag set v4.2.0

02 `--no-promo` suppresses first-run audit prompt

NEW

80

Adds `--no-promo` and `IFIXAI_NO_PROMO=1` to suppress the new interactive audit prompt on first run. The prompt is never shown in CI, non-TTY environments, or on `--dry-run` , keeping scripted and agent-driven runs unblocked.

Prevent the one-time audit prompt from blocking a CI pipeline by setting the environment variable.

```
$ IFIXAI_NO_PROMO=1 ifixai run --provider openai --api-key "$OPENAI_API_KEY" --model gpt-4o
```

– Flag, env var, and a runnable CI example all given v4.1.0

03 Vercel and Cloudflare AI Gateway provider support

NEW

77

Adds `--provider vercel` and `--provider cloudflare` to `ifixai run` for routing inspections through Vercel AI Gateway and Cloudflare AI Gateway respectively.

Route an inspection run through Vercel AI Gateway instead of hitting the provider directly — useful for unified logging and rate-limit management.

```
$ ifixai run --provider vercel --api-key "$VERCEL_AI_GATEWAY_KEY"
--model openai/gpt-4o
```

– Exact flags plus a runnable example command v4.2.0

04 Grading accuracy fixes and stricter resume/min-score checks

 75

BREAKING

Judge-verdict and grading fixes mean runs previously over-credited under v4.0.0 will score lower or show more unscored probes. An invalid `--min-score` value (NaN, Inf, or outside 0-1) is now a hard usage error that aborts the run before it starts. A `--resume` run whose endpoint or system prompt has changed is now refused rather than proceeding with mismatched results.

– Flags and failure conditions named but no example run v4.1.0

05 'ifixai trajectory' command for policy-conformance checks

NEW

 70

Checks a recorded sequence of agent actions against a business policy, catching double-used approvals, out-of-scope delegations, overspent budgets, and wrong citations. Reads JSONL or OTLP and supports offline replay.

– Command and input formats named but no runnable example given v4.2.0

06 Stricter CLI exit codes and startup validation

BREAKING

 68

CI exit codes changed: RUN INVALID now exits 2 regardless of score, a partial run exits 1, and a scoped run without `B01`, `B08` or `P01` exits 0 only if every selected inspection passed, otherwise exits 2. Configuration values that were previously accepted silently are now hard errors at startup. Ctrl-C at an interactive prompt now aborts instead of accepting the default.

– Precise exit-code semantics but no example invocation v4.2.0

Finds harmful action sequences composed of individually permitted tools, then suggests the smallest permission removal that keeps legitimate goals reachable. Runs entirely offline.

– Mechanism described but no example command or flags shown v4.2.0

↳ BREAKING ON UPGRADE

- ! `ifixai run` now runs the 10-inspection `essential` suite by default instead of all 60; pipelines relying on full coverage must add `--full-run`.
- ! CI exit codes changed: RUN INVALID now exits 2 regardless of score; a partial run exits 1; a scoped run without B01, B08 or P01 exits 0 only if every selected inspection passed, otherwise exits 2.
- ! Configuration values that were previously accepted silently are now hard errors at startup.
- ! Ctrl-C at an interactive prompt now aborts instead of accepting the default.
- ! Runs previously over-credited under v4.0.0 will score lower or show more unscored probes after the judge-verdict and grading fixes.
- ! An invalid `--min-score` value (NaN, Inf, or outside 0-1) is now a hard usage error that aborts the run before it starts.
- ! A `--resume` run whose endpoint or system prompt has changed is now refused rather than proceeding with mismatched results.

WAS THIS USEFUL?



Straiker

Feedback



Straiker Brings Independent AI Agent Red-Teaming to Workday Agent Passport (2026-10-16)

BLOG ↗

Straiker discovers and inventories enterprise AI agents, tests them with automated red teaming and protects them with runtime guardrails.

Ascend AI integrates with Workday Agent Passport to issue cryptographically signed, version-pinned Stamps evidencing independent adversarial red-teaming of AI agents.

- › Ascend AI submits signed assessment results to Workday Agent Passport through Workday's partner interfaces, surfacing as cryptographically signed Stamps in the Agent System of Record.
- › Adds an **Adversarial Resilience** Trust Category to Workday Agent Passport, with version-pinned Stamps that are automatically invalidated when the tested agent changes.
- › Workday Marketplace customers can configure which Stamps agents must hold and inspect Stamp details — risk tested, tester identity, agent version — before installation.

COVERAGE

NEW detection [Adversarial Resilience testing for AI agents via Ascend AI](#)

NEW integration [Straiker Ascend AI integration with Workday Agent Passport](#)

[All of Straiker's coverage changes →](#)

WAS THIS USEFUL?



code-on-incus

[Feedback](#)[v0.15.0](#)[NOTES](#)[RANK](#)

Give the agent a machine. Just not yours.

coi v0.15.0 ships a full CLI for running AI coding agents (Claude Code, Codex, opencode, pi, omp) in isolated Incus containers, adding session execution, kernel-level threat monitoring, named profiles, container lifecycle management, network/tool configuration, and supply-chain guards.

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

01 Core sandbox execution commands NEW 95

`coi shell` drops into an interactive AI coding session inside a full-OS Incus container, with the project mounted at `/workspace` and host credentials excluded by default. `coi run -- <cmd>` executes any command in the sandbox with streamed output and propagated exit codes, and `coi run --prompt-name <name>` runs headless agents from predefined prompts. `coi build` (with a `coi build sandbox` variant) constructs the base container image.

– Names exact commands, flags and mount path; fully actionable. [v0.15.0](#)

02 Kernel-level threat monitoring and diagnostics commands NEW 95

`coi monitor` provides a real-time security dashboard with kernel-level detection of reverse shells, C2, exfiltration, and DNS tunneling, auto-pausing sessions on HIGH severity and auto-killing on CRITICAL severity. `coi audit` streams the JSONL threat-event log for piping into a SIEM or `jq`. `coi top` reports per-container CPU, memory, and IO metrics resolved to workspace and alias. `coi health` diagnoses the coi setup.

– Specific commands and severity behavior, all runnable. [v0.15.0](#)

03 Named profile management for containers NEW 95

`coi profile create <name>` and `coi profile list` scaffold and manage reusable named container profiles stored as `~/.coi/profiles/<name>/config.toml`. `coi schema profile` and `coi validate profile` provide JSON-schema validation of profile configs. The `--profile <name>` flag on `coi shell` applies a named profile (image, tool, limits, network mode, build scripts, agent instructions) in one flag, with `inherits` support to build on another profile.

Lock an untrusted agent session to restricted egress and 4 CPU cores using a named profile — without touching global config.

```
$ coi shell --profile hardened
```

– Exact commands, config path and a runnable example given. v0.15.0

04 Network and tool configuration in coi.toml NEW 95

A `[network] mode` config key supports `open`, `restricted`, and `allowlist` values backed by nftables egress control and DNS pinning. A `[tool] name` config key accepts `claude`, `codex`, `opencode`, `pi`, and `omp`, with `permission_mode` set to `bypass` or `interactive`, read from `~/.coi/config.toml` or `./.coi/config.toml`. Per-repo `.coi/config.toml` files share image, network mode, and limits across a team.

Drop a per-repo config so every team member gets the same image, network mode, and tool when they run `coi shell` from the project root.

```
toml

[tool]
name = "claude"
permission_mode = "bypass"

[network]
mode = "restricted"

[container]
image = "coi-default"
persistent = true
```

– Exact config keys, values and file paths with a worked example. v0.15.0

05 Container lifecycle management commands NEW 85

`coi list --all` shows active containers and saved sessions, `coi shutdown` and `coi kill` stop or force-kill containers, `coi clean` removes stopped containers and orphaned resources, and `coi attach` reconnects to a running session.

– Named commands with clear purpose but little mechanism detail. v0.15.0

THINNER COVERAGE BELOW

06 Supply-chain guard protections NEW 55

Git hooks and IDE configs are mounted read-only, protected paths are enforced, and a branch guard blocks direct commits and pushes to `main`.

– Describes protections but no named config or command to invoke. v0.15.0

WAS THIS USEFUL?



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

h5i's v0.5.0 and v0.4.9 releases bundle the browser rendering engine directly into the `h5i` binary, introduce an installable plugin system covering HTTP security testing, endpoint discovery, and CI security regression testing, and ship an Axum-based Rust framework (`h5i-app`) for formally verifying application logic in Lean 4.

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

01 Security regression testing plugin (test) NEW 95

`h5i plugin install test --from <path>` installs `h5i-test`, which replays repository-owned attack flows and writes results as JSON and JUnit. `--min-coverage` on `h5i test` enforces OpenAPI coverage thresholds (report-only without it), and `--in` on `h5i test` records a witness outside the browser for each request on a sandboxed host.

Run security regression tests against a local app in CI, failing the build if OpenAPI coverage drops below 80%.

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

Install the security regression test plugin and run saved attack flows against a local dev server, failing CI if OpenAPI coverage drops below a threshold.

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

– Flags, output formats, and runnable CI command all given.

v0.5.0

v0.4.9

02 Lean 4 formal verification framework (h5i-app) NEW 95

The `h5i-app` Axum-based Rust framework, added via `h5i-app = { version = "0.1.1", features = ["http", "postgres"] }` in `Cargo.toml`, provides `h5i app new`, `h5i app extract`, `h5i app prove`, and `h5i app mutate` subcommands to translate application logic into Lean 4 and formally verify it.

Add h5i-app to a Rust project to enable Lean 4 formal verification of authorization and business logic.

```
toml

h5i-app = { version = "0.1.1", features = ["http", "postgres"] }
```

Bootstrap a new verifiable Rust web app, translate its logic to Lean 4, and run formal proofs against it.

```
h5i app new counter && cd counter
h5i app extract
h5i app prove
```

Add `h5i-app` as a dependency to start building a formally verifiable Axum service.

```
toml

[dependencies]
h5i-app = { version = "0.1.1", features = ["http", "postgres"] }
```

– Dependency snippet and full command sequence make it runnable.

v0.5.0

v0.4.9

03 Browser engine bundled into h5i binary

IMPROVED

75

The browser rendering engine is now bundled directly into the `h5i` binary instead of the previously separate `h5i-browser-light` archive, exposed via `h5i browser open`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` subcommands, plus `h5i __engine` for direct CLI access to the embedded engine.

– Names all subcommands and prior archive, but no worked example.

v0.5.0

v0.4.9

04 Recon plugin for endpoint discovery

NEW

75

`h5i plugin install recon --from <path>` installs the `h5i-recon` discovery plugin, which records exposed endpoints and confirms candidates only when a live request answers for them.

Install the recon plugin after downloading its archive, then crawl a target to discover endpoints and attack surfaces.

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

– Confirmation mechanism explained with a runnable crawl command.

v0.5.0

05 **HTTP security workbench plugin (websec)** NEW 70

`h5i plugin install websec --from <path>` installs the `h5i-websec` HTTP security workbench, offering request capture, replay, and manipulation via `h5i websec`; it runs browser verbs through the same policy, budget, and receipts as manual sessions.

– Mechanism and install command named, but no usage example shown. v0.5.0 v0.4.9

THINNER COVERAGE BELOW

06 **Static builds for Linux, macOS, and Windows** NEW 50

Supports static musl builds for Linux x86_64 (`x86_64-unknown-linux-musl`) and Linux arm64 (`aarch64-unknown-linux-musl`), plus macOS Apple Silicon (`aarch64-apple-darwin`) and Windows x86_64 (`x86_64-pc-windows-msvc`).

– Exact target triples named but no install steps given. v0.5.0

07 **Local dashboard via `h5i ui`** NEW 45

Adds `h5i ui` command to open a local dashboard for monitoring browser and session activity.

– Command named but dashboard contents not detailed. v0.5.0

08 **`h5i plugin list` command** NEW 40

Adds `h5i plugin list` to display currently installed plugins.

– Bare command name, no further behavior described. v0.5.0

WAS THIS USEFUL?



pipelock

[Feedback](#)

v3.7.0

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

Open-source AI agent firewall for MCP security and agent egress. Scans mediated HTTP, MCP, A2A, and WebSocket traffic for exfiltration, SSRF, and prompt injection, and emits mediator-signed action receipts: verifiable audit evidence from outside the agent.

pipelock v3.7.0 adds parallel signed receipt chains with group verification, enforced filesystem containment with canary-checked grants, a `pipelock exec` command, and kernel-observed MCP service identities, alongside a wide set of breaking changes tightening chain-verification exit codes, credential protection, schema canonicalization, and proxy header handling.

←→ WHAT SHIPPED • 19 FEATURES most completely described first [? what's the number?](#)

01 Parallel signed receipt chains in flight recorder NEW

95

Adds `flight_recorder.receipt_chains` (accepts 1–32, defaults to 1) to enable parallel signed receipt chains, verified together via `pipelock chain --dir --group ID` or `pipelock group --group-id ID` using TypeScript/Rust verifiers 0.6.5 (`@pipelock/verifier-ts`, `pipelock-verifier-rs`); `pipelock-verify` 0.4.1 on PyPI verifies individual shards but not the whole group. Changing `receipt_chains` requires a restart, multi-chain automatic anchoring is unsupported, and each native evidence run has a 256 MiB verifier ceiling — a restart over a run that reached it is refused when `receipt_chains` is above 1.

Enable eight parallel receipt chains on a high-throughput recording proxy to distribute signing load across shards.

```
yaml

flight_recorder:
  receipt_chains: 8
```

Verify all eight signed receipt chain shards together after a recording run to confirm a complete group.

```
$ pipelock chain --dir /var/lib/pipelock/receipts --group my-
group-id
```

— Names exact config key, commands, verifier versions, and limits.

v3.7.0

02 Filesystem containment enforcement in agent launches NEW

Adds `containment.filesystem.mode: enforce` to hide other home directories and signing material, mount unexpired workspace grants, and refuse launch unless a canary proves the containment profile took effect; this mode defaults to `enforce` on new installs, while existing configs that omit it remain off. Adopting the related socket isolation requires running `contain install` with the new binary and restarting agent services and sessions — existing managed units and wrappers are not updated automatically.

Enforce filesystem containment on a new install to hide host home directories and signing material from proxied workloads.

```
containment:
  filesystem:
    mode: enforce
```

– Exact config key, default behavior, and migration command given. v3.7.0

03 Writable state directory requirement for issuer-cookie evidence

86

BREAKING

Issuer-cookie evidence requires a writable `$XDG_STATE_HOME` (default `~/.local/state`); the Helm chart and `pipelock init sidecar` output use a read-only root filesystem, causing write failures and evidence resets. Set `XDG_STATE_HOME` to a writable directory using `extraEnv`, `extraVolumes`, and `extraVolumeMounts` with the chart.

– Names the env var, default path, and exact Helm remediation keys. v3.7.0

04 Exit code change for incomplete chain verification

BREAKING

78

Chain verification now exits 3 and reports `VALID PREFIX, END UNPROVEN` for an open prefix instead of exit 0; automation treating exit 0 as completeness must be updated. Use `--allow-prefix` only when deliberately checking an open prefix, and verify recorder directories with no end marker using `--whole-recorder`.

– Exact exit code and flags given but no worked example. v3.7.0

05 Stop budget and upgrade path for recording proxies

BREAKING

75

Recording proxies now have a 10-minute stop budget; hand-written deployments must set their own termination grace period. `pipelock contain upgrade` from 3.7.0 adds it to existing units unless one is already set, while upgrading from an earlier release requires running `pipelock contain install` with the new binary.

– Gives exact budget and the upgrade commands to run. v3.7.0

06 Acknowledged findings for MCP tool credential requests NEW

71

Adds `mcp_tool_scanning.acknowledged_findings` to bind an operator-reviewed MCP tool credential-request finding to the configured server, tool definition, detector revision, and matches via a deployment-held `__HMAC__` key.

– Named config key and mechanism, but no CLI example. v3.7.0

07 Kernel-observed MCP service identities NEW

70

Adds `mcp_identities` config and the `pipelock mcp identity` subcommand to register and inspect local MCP services with kernel-observed identities, checking service owner and pinned program inputs on each new connection.

– Named config and subcommand, no worked example. v3.7.0

08 Canonical key enforcement for schema-bound JSON

BREAKING

68

Schema-bound JSON (typed MCP, `A2A`, GraphQL, contract, and receipt inputs) now rejects noncanonical aliases and case-variant duplicate keys; legacy signed receipts with noncanonical-case keys are rejected even when their signatures verify. Replay profiles v4 and v5 pin byte decoding, while historical profiles retain their prior contracts.

– Names affected input types and replay profile versions. v3.7.0

09 Correlation header propagation across event sinks NEW

65

Adds `emit.correlation_header` to copy a bounded, `DLP`-clean request header into webhook, syslog, `CEF`, `OCSEF`, and `OTLP` events on HTTP-capable transports.

– Names config key and five target formats, no example. v3.7.0

10 Per-request TLS-interception timing NEW

65

Adds per-request TLS-interception timing to `intercept_http`, recording status, response size, total duration, upstream-header wait, and client cancellation, honoring `logging.include_allowed`.

– Names fields recorded and the config flag honored. v3.7.0

11 Bounded limits on receipt scanning reassembly

Receipt scanning reassembly now has explicit bounded limits: splits requiring multiple fixed values, dropped enum text, nonadjacent decision-text fragments, five or more pieces, four pieces outside forward/reverse order, or three pieces on receipts too wide for that search are no longer checked. Review the receipt content boundary before relying on throughput improvements.

– Lists exact limit conditions but no config surface to tune them. v3.7.0

12 pipelock exec for proxy-aware command launches NEW

64

Adds `pipelock exec -- COMMAND` to check a running proxy's health, forwarding, and kill switch, then launch cooperative clients with shared proxy and CA environment settings.

– Names the subcommand but no worked example run. v3.7.0

13 Child environment override restrictions for MCP/sandbox launches

60

BREAKING

Explicit child environment overrides that select loaded code, such as `PYTHONHOME`, `NODE_PATH`, and `CLASSPATH`, are now refused for MCP and sandbox launches.

– Names exact blocked environment variables. v3.7.0

THINNER COVERAGE BELOW

14 Proxies reject partial origin replies on Range requests

59

BREAKING

Forward, intercepted, and reverse proxies now strip Range and cache-validator request headers and refuse origin `206` or `304` replies as `response_incomplete`; Browser Shield exemptions do not bypass this.

– Names the status codes and error label, no config toggle. v3.7.0

15 Destination grants for exact TCP endpoints NEW

58

Adds `destination_grants` to authorize an exact TCP host and port through existing destination checks without disabling content scanning or protected-address checks.

– Names the config key and scope, no example given. v3.7.0

16 DLP action enforcement for credential protection

BREAKING

58

Core credential protection now blocks in audit mode instead of only warning, and JWT cookies now follow the configured DLP action and core floor instead of always warning; a fully redacted core credential on HTTP request bodies forwards only to a configured request-body trusted host, in both audit and enforcement modes.

– Describes new blocking behavior and one carve-out, no config key named. One-line change with no named config key or example.

v3.7.0

17 UUIDv7 request identifiers

BREAKING

55

Generated request IDs use UUIDv7 instead of a counter; consumers of logs, events, and receipts must accept the `req-<UUIDv7>` format.

– Names the exact new ID format for downstream consumers.

v3.7.0

18 Canonical policy hash change for credential-audience rules

BREAKING

55

Built-in credential-audience changes move the canonical policy hash; operators must update conductor and followers together before republishing fleet policy and refresh any consumers that pin the prior default hash.

– Gives migration sequence but no concrete hash or config name.

v3.7.0

19 Signed recovery seals for torn recordings

NEW

50

Adds signed recovery seals so a fresh run can attest the exact torn predecessor bytes and the new opening, with verifiers reporting an attested discontinuity while preserving the damaged history's failure.

– Explains mechanism but no named config or command.

v3.7.0

↳ BREAKING ON UPGRADE

- ! Chain verification now exits 3 and reports `VALID PREFIX, END UNPROVEN` for an open prefix instead of exit 0; automation treating exit 0 as completeness must be updated. Use `--allow-prefix` only when deliberately checking an open prefix. Recorder directories with no end marker must be verified with `--whole-recorder`.
- ! Receipt scanning reassembly now has explicit bounded limits; splits requiring multiple fixed values, dropped enum text, nonadjacent decision-text fragments, five or more pieces, four pieces outside forward/reverse order, or three pieces on receipts too wide for that search are no longer checked. Review the receipt content boundary before relying on throughput improvements.
- ! Run `contain install` with the new binary and restart agent services and sessions to adopt socket isolation; existing managed units and wrappers are not automatically updated. `containment.filesystem.mode` defaults to `enforce` on new installs; existing configs that omit it remain off.

- ! Explicit child environment overrides that select loaded code (such as `PYTHONHOME` , `NODE_PATH` , and `CLASSPATH`) are now refused for MCP and sandbox launches.
- ! Schema-bound JSON (typed MCP, A2A, GraphQL, contract, and receipt inputs) now rejects noncanonical aliases and case-variant duplicate keys. Legacy signed receipts with noncanonical-case keys are rejected even when their signatures verify. Replay profiles v4 and v5 pin byte decoding; historical profiles retain their contracts.
- ! Core credential protection blocks in audit mode (previously only warned). A fully redacted core credential on HTTP request bodies forwards only to a configured request-body trusted host in both audit and enforcement modes.
- ! `flight_recorder.receipt_chains` accepts 1–32 and defaults to 1; parallel receipt groups require group-capable verifiers (`@pipelock/verifier-ts` and `pipelock-verifier-rs` 0.6.5). Changing `receipt_chains` requires restart. Multi-chain automatic anchoring is unsupported.
- ! Each native evidence run has a 256 MiB verifier ceiling; a run that reaches it cannot publish a close that verifies. With `flight_recorder.receipt_chains` above 1, a restart over such a run is refused.
- ! Recording proxies now have a 10-minute stop budget; hand-written deployments must set their own termination grace period. `contain upgrade` from 3.7.0 adds it to existing units unless one is already set; upgrading from an earlier release requires running `pipelock contain install` with the new binary.
- ! Issuer-cookie evidence requires a writable `$XDG_STATE_HOME` (default `~/.local/state`); the Helm chart and `pipelock init sidecar` output use a read-only root filesystem, causing write failures and evidence resets. Set `XDG_STATE_HOME` to a writable directory using `extraEnv` , `extraVolumes` , and `extraVolumeMounts` with the chart.
- ! Forward, intercepted, and reverse proxies now strip Range and cache-validator request headers and refuse origin `206` or `304` replies as `response_incomplete` ; Browser Shield exemptions do not bypass this.
- ! JWT cookies now follow the configured DLP action and core floor instead of always warning.
- ! Generated request IDs use UUIDv7 instead of a counter; consumers of logs, events, and receipts must accept `req-<UUIDv7>` .
- ! Built-in credential-audience changes move the canonical policy hash; update conductor and followers together before republishing fleet policy and refresh consumers that pin the prior default hash.

WAS THIS USEFUL?



PLATFORMS & UTILITIES

Microsoft Sysinternals

Feedback



ProcExp, ZoomIt, BgInfo, ListDLLs, ProcDump, ProcMon, PsKill, Sigcheck, ZoomIt for macOS

BLOG ↗

RANK

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

Microsoft shipped a major Process Explorer 17.2 update with new memory, kernel and hardware visibility features, plus a UTF-8 output option in Sigcheck 2.93.

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

01 Process Explorer 17.2 diagnostics and visibility upgrades

NEW

51

Process Explorer v17.2 adds live kernel memory dump creation, an efficiency mode, a memory map explorer, NPU monitoring, persistent activity history, and a new Devices page added to the System Information view.

– Names several new views/features but no mechanism or exact UI steps

ProcExp, ZoomIt, BgInfo, ListDLLs, Proc...

02 UTF-8 output encoding in Sigcheck NEW

47

Sigcheck v2.93 adds a UTF-8 encoding option for file output.

– Names the concrete format option but no flag syntax given

ProcExp, ZoomIt, BgInfo, ListDLLs, Proc...

WAS THIS USEFUL?



// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 17, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe

Feedback on this issue