THE CYBER TOOLCHAIN

NO. 068

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 8, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 161 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 25 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. GitGuardian's AI Hooks intercept secrets on their way out through AI coding agents and MCP tool calls, ak path that pre-commit hooks never see. Firezone now gates

vim ? help

tunnels on an X.509 device certificate, not just a logged-in user. DefectDojo lets a KEV or exploit-evidence flag drive rules, so triage ranks itself.

PROTECT

Stop a secret before an AI agent writes it somewhere you don't control

GitGuardian's AI Hooks sit on the agent and MCP tool-call path rather than on `git commit`, closing the gap where an assistant pastes a live credential into a file, a prompt, or an outbound API call; `ggshield machine setup` and `machine doctor` mean rolling this to a laptop is one command plus a health check instead of per-repo hook plumbing. ZeroPath's org-wide toggle to redact AI-generated descriptions on secret findings addresses the mirror-image problem — the finding itself quoting the secret into a ticket.

GitGuardian SaaS · ZeroPath

PROTECT

Require a device certificate, not just a valid login, before a tunnel comes up

Firezone's X.509-backed device trust ties a Client session to a key the endpoint holds, so stolen SSO credentials on an unmanaged machine no longer get network access — the check happens at connect time rather than as a posture report you read later. Auth0 moving Advanced MFA Configurations to GA lets factor policy be set per app and tenant in production instead of a global lowest-common-denominator setting.

Firezone · Auth0

PROTECT

Let exploitation evidence pick the next fix, and keep the exception without killing the scanner

DefectDojo rules can now branch on KEV membership and exploit evidence, which turns "which of these 4,000 findings actually matters" from a weekly spreadsheet pass into a condition the pipeline evaluates on ingest; the inline repository fix flow shortens the hop from finding to PR. Mend's Top Fix column collapses many findings onto the one dependency bump that clears them, and its IDE/AI-agent suppression tool records a false positive where the developer is sitting rather than in a separate console.

defectdojo · Mend Platform

◆ AI-Powered Offensive Security

RAPTOR

Feedback ⓘ 1 RELEASE · 2026-09-07 NOTES ↗ RANK

RAPTOR chains static analysis, binary analysis, LLM vulnerability validation, exploit generation, and patch writing into one autonomous security research workflow.

> RAPTOR v3.1.0 adds compiler-analyzer and expanded-Semgrep scanning channels, finer sandbox audit controls, OCI/Dockerfile/tar scan substrates, a software composition analysis pipeline, multi-LLM provider support, and Frida-based reachability integration, plus a wide set of internal pipeline additions for sink discovery, checker synthesis, and web correlation.

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

01 **Compiler-analyzer scan channel** NEW 90

Adds `--compiler-scan` / `--no-compiler-scan` flags to enable or disable a compiler-analyzer channel using gcc `-fanalyzer` or clang `--analyze`, plus `--compiler-scan-max-tus <n>` (default 2000) to cap the number of translation units analyzed.

Run the compiler-analyzer channel on a C/C++ codebase alongside Semgrep, capping the number of translation units to keep the scan tractable on a large repo.

```
$ /scan /path/to/c-project --compiler-scan --compiler-scan-max-tus 500 --build-command 'make -j4'
```

— Exact flags, tool names, default value, and runnable example. v3.1.0

02 **Expanded-Semgrep macro scanning** NEW 85

Adds `--expanded-semgrep` flag to re-run Semgrep rules over preprocessor-expanded views of macro-heavy C/C++ translation units, catching bugs that only appear in expanded output and that plain Semgrep misses.

Scan a macro-heavy C/C++ codebase and catch bugs that only appear in expanded preprocessor output, which plain Semgrep misses.

```
$ /scan /path/to/c-project --expanded-semgrep --languages c,cpp
```

— Named flag, clear mechanism, and runnable example. v3.1.0

03 **Sandbox audit mode verbosity and budgeting** NEW 80

Adds `--audit-verbose` flag for strace-style verbose output during sandbox audit runs, and `--audit-budget <n>` flag to cap the maximum number of audit events collected per run.

Audit a target's sandbox behavior before enforcing — log every syscall that would have been blocked without actually blocking anything, with a cap to avoid log flooding.

```
$ /scan /path/to/target --audit --audit-verbose --audit-budget
500
```

— Named flags with a runnable example command. v3.1.0

THINNER COVERAGE BELOW

04 **Container image and tarball scan substrates** NEW 50

Adds OCI registry, Dockerfile, and tar as scan substrates, enabling RAPTOR to analyze container images and tarballs directly.

— Named formats but no invocation shown. v3.1.0

05 **CVE environment modeling package** NEW 45

Adds `cve-env` as an integrated package (`packages/cve_env`) for CVE environment modeling (Phase 1).

— Named package path but no usage instructions given. v3.1.0

06 **Multi-LLM provider support** NEW 45

Adds full multi-LLM provider support for the analysis dispatch layer, with Mantle and InvokeModel coexistence.

— Named providers but no configuration steps shown. v3.1.0

07 **Mechanical sink discovery in understand module** NEW 40

Adds mechanical sink discovery and framework API detection (MAP-5f) to the `understand` module.

— Named module and internal ID but no usage path. v3.1.0

08 **Web correlation tier with poisoned-render oracle** NEW 40

Adds a web correlation tier with expectations, check legs, and a poisoned-render oracle for web vulnerability analysis.

— Named components but no command or endpoint given. v3.1.0

09 **Per-function validation annotations** NEW 35

Adds per-function annotations emitted across validation stages A, B, and D via the `validate` module.

— Named module and stages but no interface given. v3.1.0

10 **Agentic checker synthesis** NEW 35

Adds agentic checker synthesis — the system now generates a checker per confirmed finding.

— Explains behaviour but no configuration or command. v3.1.0

11 **Software composition analysis pipeline** NEW 30

Adds a full SCA (software composition analysis) mechanical pipeline with a supply-chain detector and first-party hardening.

— Described only in prose, no mechanism or command given. v3.1.0

12 **JS and Go reachability framework registration** NEW 30

Adds JS and Go function-as-argument framework registration for reachability analysis.

— Named languages but no usage detail given. v3.1.0

13 **Frida evidence pipeline for reachability** NEW 30

Adds Frida evidence pipeline with full consumer wiring and reachability integration.

— Named tool integration but no usage instructions. v3.1.0

14 **Corpus runner for rule calibration** NEW 25

Adds a corpus runner with rule evaluation and label linting for calibrating detection rules.

— Named component but no command or config shown. v3.1.0

15 **Persistent verdict-bearing coverage store** NEW 25

Adds a persistent verdict-bearing coverage store wired into the live pipeline.

— Bare description of internal store, no interface shown. v3.1.0

16 **Proxy-hosts auto-calibration for CodeQL** NEW 25

Adds proxy-hosts auto-calibration for CodeQL analysis.

— Named integration but no mechanism or usage shown. v3.1.0

17 **Source intelligence substrate** NEW 20

Adds a source intelligence substrate with a minimal axis-1 rule for source-level intel gathering.

— Bare description, no mechanism or usage detail. v3.1.0

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

Invicti

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗ RANK

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

> Invicti brings Agentic Pentesting management and a new Vulnerability Chaining Agent into the AppSec Platform UI, adds rate-limit-aware scan pacing and WebRTC chatbot detection to the DAST engine, and expands scan-cli's `INVICTI_MINIMUM_CONFIDENCE` and `INVICTI_IGNORE_SSL_CHECKS` environment variables to behave more independently and broadly.

└──▷ **WHAT SHIPPED · 19 FEATURES** most completely described first ? what's the number?

01 **Broader scope for SSL check bypass** IMPROVED 90

The `INVICTI_IGNORE_SSL_CHECKS` environment variable in scan-cli now applies to all outbound platform API calls made during a scan — including DAST, Inventory, Identity, and report/export endpoints — not just the ephemeral agent.

Allow scan-cli to reach an Invicti platform instance behind a self-signed TLS certificate, ensuring all API calls (DAST, Inventory, Identity, report/export) bypass SSL verification.

```
$ export INVICTI_IGNORE_SSL_CHECKS=true
scan-cli scan --target-url https://app.internal.example.com
```

— Exact env var, scope change, and runnable example changelog-20260908-29675a96

02 **Independent confidence-based scan failure** IMPROVED 85

The `INVICTI_MINIMUM_CONFIDENCE` environment variable in scan-cli now independently fails a scan when its threshold is exceeded, without also requiring `INVICTI_MINIMUM_SEVERITY` to be set.

Fail a CI pipeline when scan confidence exceeds a threshold, without also gating on severity — useful when you want confidence-only quality gates.

```
$ export INVICTI_MINIMUM_CONFIDENCE=confirmed
# INVICTI_MINIMUM_SEVERITY is intentionally not set
scan-cli scan --target-url https://app.example.com
```

— Exact env var with runnable example command changelog-20260908-29675a96

03 **Empty severity threshold in CI/CD script generator** IMPROVED

75

Leaving the minimum severity field empty in the CI/CD script generator now produces a script with no minimum vulnerability threshold, applying to Azure Pipelines, CircleCI, GitHub Actions, GitLab CI/CD, and Jenkins.

Generate a CI/CD script in GitHub Actions with no minimum severity gate, so the pipeline reports all findings without blocking on severity.

> 📍 In the Invicti platform, go to Integrations › CI/CD Script Generator, select GitHub Actions, leave the Minimum Severity field empty, and download the generated script.

– Names all supported CI platforms with exact UI steps changelog-20260908-29675a96

04 **IDOR/BOLA authentication credential improvements** IMPROVED

70

IDOR/BOLA authentication now supports Invicti Secrets Manager credentials (API key, basic auth, and bearer token types), and the character limit for authentication values in IDOR/BOLA credentials has been increased from 1,024 to 4,096 characters.

– Exact limits and credential types with UI screenshots changelog-20260908-29675a96

05 **Recon agent document import** NEW

65

The Recon agent can now import architectural diagrams and documentation files in `PDF`, `DOCX`, and `MD` formats, enabling discovery of non-crawlable attack surfaces.

– Names input formats but not the exact import interface changelog-20260908-29675a96

06 **Attack Agents visibility and resource controls** IMPROVED

65

Assessment Profiles now support input caps and credit spend limits for more granular control over agent resource use; the OCTO sitemap now supports excluding URL paths and parameters, bringing it in line with DAST scan exclusion capabilities; and the DAST Engine now appears as a distinct entry under Attack Agents so total vulnerability counts align across AI agents and traditional DAST scans.

– Names three concrete controls without exact config keys changelog-20260908-29675a96

07 **Rate-limit-aware scan pacing** IMPROVED

65

The scan engine now automatically throttles request pacing when it receives Retry-After or `X-RateLimit-*` rate-limiting response headers, reducing incomplete coverage and false-negative findings, and applies exponential backoff with jitter when

retrying requests that receive a `429 Too Many Requests` response with no `Retry-After` header.

— Automatic behavior; no config toggle mentioned changelog-20260908-29675a96

08 **Agentic Pentesting management in AppSec Platform UI** NEW

60

Agentic Pentesting can now be managed directly from the AppSec Platform UI, with its findings flowing into the same pipeline as standard DAST scans for unified visibility and triage.

— Describes mechanism but no exact UI navigation path changelog-20260908-29675a96

09 **Bearer Token and API Key header types** NEW 60

`Bearer Token` and `API Key` are now dedicated, first-class authentication header types in target configuration, separate from functional headers, for easier authenticated scan setup.

— Names new header types but no exact config path changelog-20260908-29675a96

THINNER COVERAGE BELOW

10 **Vulnerability Chaining Agent** NEW 55

A new Vulnerability Chaining Agent correlates individual findings to execute multi-step vulnerability chains, uncovering complex attack paths that single-step analysis might miss.

— Named agent with mechanism but no invocation surface changelog-20260908-29675a96

11 **Scan schedules for applications and collections** NEW

55

Adds the ability to create scan schedules for an entire application or collection, automatically running DAST scans across all included targets on a chosen recurrence.

— Describes scope and recurrence but no config/API name changelog-20260908-29675a96

12 **SCIM2 provisioning toggle** NEW 55

Adds the ability to toggle SCIM2 provisioning endpoints on or off for an organization, giving administrators control over whether user and team updates or deletions through SCIM2-provisioned accounts are permitted.

— Clear admin control but no exact endpoint or toggle path changelog-20260908-29675a96

13 **DAST technology fingerprinting for Attack Agents** NEW

50

DAST technology fingerprinting automatically shares detected technology and framework stacks with Attack Agents so exploit payloads are tailored to the target's

actual stack.

— Clear mechanism but no configuration surface named changelog-20260908-29675a96

14 **Assessment reports in Markdown and JSON** NEW 50

Assessment reports can now be downloaded in Markdown (`.md`) and raw JSON formats.

— Names formats but no endpoint or UI path given changelog-20260908-29675a96

15 **Helm chart support for OpenShift and ROSA** NEW 50

The scalable internal scanning agent Helm chart now supports deployment on OpenShift and Red Hat OpenShift Service on AWS (ROSA) clusters.

— Names platforms but no chart values or commands changelog-20260908-29675a96

16 **Path- and input-level coverage visualization** NEW 45

The sitemap now shows path-level and input-level coverage visualization, indicating which endpoints were analyzed by scan-wide and input-specific agents.

— Explains what's shown but no navigation or config detail changelog-20260908-29675a96

17 **WebRTC chatbot detection** NEW 45

Adds WebRTC data channel (SDP) chatbot detection, expanding chatbot coverage to applications using real-time communication protocols.

— Names protocol but no configuration or scan option given changelog-20260908-29675a96

18 **API target creation from API Catalog** NEW 45

Enables creating a new API target directly from the API Catalog page, with or without an API specification attached.

— Clear UI starting point, limited mechanism detail changelog-20260908-29675a96

19 **Granular permissions in Role Management** IMPROVED 35

Granular permission operations are now visible and configurable in Role Management, giving administrators full control over custom role definitions.

— Vague on which permissions or exact UI location changelog-20260908-29675a96

WAS THIS USEFUL?  

DataDome

Feedback ⓘ 1 RELEASE · 2026-09-07 BLOG ↗ RANK

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

> DataDome's CloudFront v3 upgrade splits exclusion configuration into dedicated environment variables, adds new default URI inspection settings, and migrates the Lambda@Edge module to ES Modules with async/await for Node.js 24 compatibility, while the Nginx module extends static-asset exclusions and adopts API-driven status codes.

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

01 New default URI inspection settings in CloudFront v3 NEW

85

Two new request inspection settings, enabled by default: `DATADOME_REMOVE_MATRIX_PARAMS` strips matrix params (`;`) from URIs before regex evaluation, and `DATADOME_ENCODED_CHAR_INCLUSION` inspects static-looking paths containing encoded characters `%23` , `%25` , `%2f` , `%3b` , or `%3f` .

Enable encoded-character inspection to catch bots probing static-looking paths with encoded chars like `%2f` or `%3f` .

```
$ DATADOME_ENCODED_CHAR_INCLUSION=true
DATADOME_REMOVE_MATRIX_PARAMS=true
```

— Both settings and their exact mechanisms are named. product docs

02 Exclusion config split into two env vars BREAKING

80

CloudFront v3 introduces `DATADOME_STATIC_ASSET_EXCLUSION` (enabled by default) to hold static-asset extension exclusions separately from custom URL exclusions, and moves custom URL exclusions to `DATADOME_URI_REGEX_EXCLUSION` exclusively. Static-asset extension patterns must be removed from `DATADOME_URI_REGEX_EXCLUSION` or they will duplicate or conflict with the new variable.

Separate your custom URL exclusions from static-asset exclusions so DataDome correctly applies both independently.

```
$ DATADOME_URI_REGEX_EXCLUSION='/health|/heartbeat'
DATADOME_STATIC_ASSET_EXCLUSION=true
```

– Names both env vars and the exact migration step required. product docs

03 **ESM and async/await migration for Lambda@Edge module** BREAKING

80

The DataDome CloudFront Node.js module migrates from CommonJS (`module.exports` / `require`) to ES Module format (`export` / `import`), requiring `.js` files to be renamed to `.mjs` . The Lambda@Edge handler moves from a callback-based signature (`exports.handler = (event, context, callback) => {}`) to `async` / `await` (`exports.handler = async (event, context) => {}`), enabling compatibility with Node.js runtime 24 and later.

– Exact before/after code signatures given but no single runnable command. product docs

THINNER COVERAGE BELOW

04 **Static asset exclusion extended to tif, tiff, pdf** IMPROVED

50

Nginx Module 3.1.0 extends the static asset exclusion regex pattern to cover `tif` , `tiff` , and `pdf` file extensions, skipping DataDome checks for those requests.

– Names the file types but no config surface to act on. Nginx Module 3.1.0 (2026-09-07)

05 **Dynamic status codes from Protection API** IMPROVED 40

Nginx Module 3.1.0 replaces hard-coded HTTP status codes with values sourced from DataDome Protection API response headers, enabling dynamic status code handling for upcoming features.

– Describes the change but gives no concrete codes or config to use.

Nginx Module 3.1.0 (2026-09-07)

└──▷ **BREAKING ON UPGRADE**

! Static-asset extension patterns must be removed from `DATADOME_URI_REGEX_EXCLUSION` ; leaving them there duplicates or conflicts with the new `DATADOME_STATIC_ASSET_EXCLUSION` variable.

! CommonJS `.js` files in the Lambda@Edge function must be renamed to `.mjs` and updated to use `export` / `import` syntax instead of `module.exports` / `require` .

! Callback-based Lambda@Edge handler signature (`exports.handler = (event, context, callback) => {}`) must be replaced with the `async` / `await` pattern (`exports.handler = async (event, context) => {}`).

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Self-Hosted

Feedback ⓘ 1 RELEASE · 2026-09-08

CHANGELOG ↗ RANK

GitGuardian Self-Hosted is an on-premise secret detection tool that scans repositories and prevents sensitive credentials from being committed.

> GitGuardian Self-Hosted sped up validity check and incident processing to real time, and extended NHI Governance identity mapping to link AWS access keys stored in secrets managers, Vault, or Kubernetes to their IAM identity.

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

01 **AWS IAM identity linking for stored secrets** NEW 50

AWS access keys stored in a secrets manager, Vault, or Kubernetes are now linked to their AWS IAM identity in the NHI Governance identity map.

– Names storage surfaces and feature area but no config steps changelog-20260908-17649a68

02 **Real-time validity check and incident processing** IMPROVED 35

Validity check results and incident creation events are now processed immediately on self-hosted instances instead of being queued, reducing the delay between a secret being flagged and an incident being created.

– Explains the change but no mechanism or metrics given changelog-20260908-17649a68

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ 1 RELEASE · 2026-09-07 NOTES ↗ RANK

Infisical is the open-source platform for secrets, certificates, and privileged access management.

> Infisical v0.165.7 adds break-glass emergency access for privileged access management requests and lets a single LDAP connection back multiple directory syncs.

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

01 **Multi-sync support for single LDAP connection** IMPROVED

45

A single LDAP connection can now back multiple Linux and Windows syncs simultaneously, reducing connection overhead for directory-integrated environments.

— States mechanism and benefit but no concrete config surface named. v0.165.7

02 **Break-glass emergency access for PAM requests** NEW

40

Adds a break-glass emergency access capability for PAM access requests, enabling out-of-band approval bypass in critical situations where normal approval flow would block access.

— Describes capability and purpose but no config/flag names given. v0.165.7

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS

Feedback ⓘ

2 RELEASES · 2026-09-07 → 2026-09-08 CHANGELOG ↗ RANK

GitGuardian SaaS is a cloud-based security platform that detects and prevents secrets leakage in git repositories and source code.

> GitGuardian shipped AI Hooks that block secrets from leaking through AI coding agents and MCP tool calls, backed by new `ggshield machine setup` / `machine doctor` commands for one-step endpoint protection, alongside continuous Docker Hub monitoring, intelligent attribution beyond the monitored perimeter, six new detectors, and expanded NHI inventory and governance capabilities.

└──▷ **WHAT SHIPPED · 18 FEATURES** most completely described first ? what's the number?

01 ggshield machine setup and machine doctor commands NEW

95

`ggshield machine setup` (requires ggshield 1.53.0+) configures AI Hooks for every detected AI agent, pre-commit/pre-push git hooks in the global git configuration, and a honeytoken decoy AWS credential profile in one idempotent run. Flags `--no-honeytokens`, `--no-ai-hooks`, and `--no-git-hooks` selectively skip individual protections. `ggshield machine doctor` reports per-AI-agent hook installation status, git hook status, and token scope validity. Planting honeytokens requires the `honeytokens:write` scope on `ggshield auth login`, gated behind the workspace Manager role, while the install script provisions default `honeytokens:check`, `endpoints:send`, and `ai-discover:send` scopes.

Bootstrap AI agent secret blocking across an entire developer machine — covers every detected agent (Codex, Cursor, Claude Code, etc.) plus git hooks in one step.

```
$ ggshield machine setup
```

Set up all three endpoint protections (AI Hooks, git hooks, honeytoken) on a developer machine in one step, including the extra scope needed to plant the decoy credential.

```
$ ggshield auth logout
ggshield auth login --scopes honeytokens:write
ggshield machine setup
```

Deploy endpoint protections across a fleet without planting honeytokens, using only the default install-script scopes.

```
$ ggshield machine setup --no-honeytokens
```

Verify that AI Hooks, git hooks, and token scopes are all correctly configured after running machine setup.

```
$ ggshield machine doctor
```

Bootstrap AI Hook and git hook protection across all detected AI agents on a developer machine in one shot.

```
$ ggshield machine setup
```

— Exact commands, flags, and required scopes given changelog-20260908-3f7375e6

See all your AI agents across the fleet…

02 **AI Hooks blocking secrets in AI agents and MCP calls** NEW

75

Blocks secrets at prompt submission, pre-tool use, and post-tool use across AI coding agents, now including Codex and Mistral Vibe alongside existing Claude Code, Cursor, Copilot CLI, and VS Code coverage; blocks surface a named secret and remediation hint. Self-hosted instances require `ClickHouse` to support machine scans with secrets findings, MCP tool calls, and full AI agent session recording.

— Names blocking points and agents but no standalone enable command

changelog-20260908-3f7375e6 See all your AI agents across the fleet…

03 **NHI inventory filtering and credential grouping** IMPROVED

75

Adds `detector`, `detector type`, `secret family`, `secret category`, and `secret provider` filters and sort options to the NHI inventory. Groups multi-match credentials (e.g. PostgreSQL username, password, host, and port) as a single inventory entry rather than separate fields, and labels unmatched secrets as `Unknown` so they remain visible and filterable rather than dropped.

— Names exact filters and behavior, UI-only action changelog-20260908-3f7375e6

04 **MCP transcript replay for post-compromise investigation** NEW

65

`ggshield ai discover --history` parses MCP tool calls from agent transcripts to reconstruct which tools an agent called, for blast-radius analysis after a suspected

credential compromise.

After a credential compromise, reconstruct which MCP tool calls an agent made by replaying its transcripts — useful for blast-radius analysis.

```
$ ggshield ai discover --history
```

After a suspected compromise, reconstruct which MCP tools an agent called by replaying its transcript history.

```
$ ggshield ai discover --history
```

— Command given but investigation output not detailed changelog-20260908-3f7375e6

See all your AI agents across the fleet…

05 **New detectors and analyzer for Firebase, Coze, Keycloak, Fish Audio, Azure, Apify, Docker Hub**

NEW 65

Adds new Firebase ID Token, Coze Access Token, Keycloak Access Token, Keycloak Refresh Token, Fish Audio API Key, and Azure Custom Vision Training Key detectors, a new Apify API tokens analyzer, and `docker login --username` / `-u` detection to the Docker Credentials detector for Docker Hub personal access tokens.

— Names every new detector, no enablement action needed product docs

06 **Vault secret classification via ggscout** NEW 60

Runs the Secrets Detection Engine on every value ggscout collects, classifying vaulted secrets against 550+ detectors including multi-match detectors (ggshield 1.54.0, ggscout 0.32.0).

— Explains mechanism but no user-facing action changelog-20260908-3f7375e6

07 **Detector accuracy improvements for New Relic, Scaleway, and Fastly**

IMPROVED 60

Extends New Relic Key and New Relic Insights Insert Key checkers to identify revoked keys via an empty POST request. Improves `scaleway_token` precision by rejecting node UUIDs from proxy-subscription configs, and improves `fastly_personal_token` precision by rejecting proxy-subscription node passwords and 32-char prefixes of longer base64url values.

— Names exact detectors and logic, no user action product docs

08 **AI agents and MCP inventory in Developer Endpoint Protection**

NEW

60

Adds AI agents and MCP inventory to Developer Endpoint Protection, reporting agents and MCP servers per endpoint, hook coverage, and whether an agent subscription is personal or company-owned.

— Explains inventory contents, no exact navigation given

See all your AI agents across the fleet…

09 **AWS access key to IAM identity linking in NHI Governance** NEW

60

AWS access keys stored in a secret manager or Kubernetes now link to their IAM identity by `AccessKeyId`, showing the permissions of the stored key in the NHI Governance exploration map.

— Names field and effect, no action step given See all your AI agents across the fleet…

THINNER COVERAGE BELOW

10 **Agent analysis coverage and overview for public incidents**

IMPROVED

55

Adds an agents overview page under Public Monitoring showing how many public incidents the agents analysed and how company-attribution verdicts break down, with one-click filtering to matching incidents. Company-attribution verdicts and risk scores are now available immediately when a public incident is created, and full agent analysis coverage is ensured on every incident, including those previously skipped.

Review how many of your public incidents the agents have analysed and drill into company-attribution verdicts.

> 📍 In the GitGuardian console, go to Public Monitoring › GitGuardian agents to see the breakdown of agent-analysed incidents and their company-attribution verdicts, then click through to the matching incidents.

— UI path given, attribution mechanism not detailed product docs

11 **Intelligent attribution beyond the monitored perimeter** NEW

55

Adds intelligent attribution beyond the monitored perimeter for new GitHub commits, GitHub commit history, and Docker Hub images — attributing public leaks to your company even when no perimeter rule matches, including surfacing past leaks that were never previously attributed.

— Names three sources but no control to configure it product docs

12 **Unified Exclusion Rules for Monitoring** IMPROVED

55

Consolidates source exclusions, file-path exclusions, and pattern exclusions into a single unified experience, managed under Settings > Monitoring > Exclusion Rules.

Consolidate all your monitoring exclusions — sources, file paths, and patterns — into one place to reduce management overhead.

> 📍 In the GitGuardian console, go to Settings › Monitoring › Exclusion Rules to add or migrate source exclusions, file-path exclusions, and pattern exclusions into the Unified Exclusion Rules experience.

— Exact UI path given, underlying mechanism thin `product docs`

---

13 **Per-team incident permission controls for Member role** NEW

50

Adds per-team incident permission controls when promoting a user to Member, including bulk application across teams and post-promotion editing of team rights.

— Describes scope but no exact UI path given `changelog-20260908-3f7375e6`

---

14 **Continuous Docker Hub monitoring** NEW

50

Adds continuous Docker Hub monitoring, scanning public images daily and attributing exposures to your company as they appear after the initial historical scan.

— Describes cadence but no configuration surface `product docs`

---

15 **Perimeter-scoped historical scans with recurring schedules**

IMPROVED

50

Scopes historical scans to each developer's time in perimeter, excluding activity from before they joined, with support for recurring scan schedules that apply new detectors to past commits.

— Explains scoping mechanism but no command or config named `product docs`

---

16 **Automatic removal of remotely-deleted sources from perimeters**

NEW

50

Adds a new workspace setting that automatically removes remotely-deleted sources from team perimeters while keeping them accessible to workspace managers.

— Names the setting's effect, not its exact location

`See all your AI agents across the fleet…`

---

17 **Multiple assignees and auto-assign leaker on incidents** NEW

35

Enables assigning multiple users to incidents and automatically assigns the leaker as assignee.

— Describes behavior but no configuration detail product docs

---

18 **Decoded JWT and Base64 payload exposure in incidents** NEW

25

Adds decoded JWT and Base64 payload exposure in incident views.

— Bare one-line description product docs

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Veracode

Feedback ⓘ 2 RELEASES · 2026-09-07 → 2026-09-08 BLOG ↗

RANK

Veracode is a cloud-based application security platform that performs static, dynamic, and software composition analysis to identify vulnerabilities in code.

> Veracode's headline work this window is operational rather than detection-side: new `veracode.yml` pipeline configuration options, a Java API Wrapper update adding pagination, retry logic, and a Findings API, and a broad expansion of Reporting API fields, filters, and a new Audit Log report. The window also adds container/IaC scan visibility and Manual Penetration Testing flaw matching on the Platform, native GitHub integration install and scan scheduling, and two breaking changes — removal of the `/healthcheck/status` endpoint and the Azure DevOps Veracode Flaw Importer task rename.

└──▷ WHAT SHIPPED · 23 FEATURES most completely described first ? what's the number?

01 New `veracode.yml` pipeline configuration options NEW

90

Adds `debug` configuration option in `veracode.yml` for debugging and troubleshooting pipeline runs (GitHub Workflow Integration v1.0.6, Azure DevOps Workflow Integration v2.3.6); adds `analysis_branch` support so scheduled scans run against a specified branch instead of the default branch (Azure DevOps Workflow Integration v2.3.10); adds a custom agent pool name parameter to specify the agent pool used for pipeline execution (Azure DevOps Workflow Integration v2.4.3); and adds support for custom predependency commands to prepare projects before scanning (GitHub Workflow Integration v2.1.3).

Enable verbose debug output in a GitHub or Azure DevOps pipeline to diagnose scan failures without modifying pipeline code.

```yaml
# In veracode.yml
debug: true
```

Pin scheduled ADO scans to a specific long-lived branch rather than the repo default, useful for release-branch gating.

```yaml
# In veracode.yml
analysis_branch: release/2026-q3
```

— Exact config keys with runnable example values Integration updates (2026-09-07)

02 **Java API Wrapper pagination, retry, and Findings API support**

NEW 85

Adds page size support to the Platform Applications API (`/appsec/v1/applications`) to improve efficiency retrieving large application lists (v26.3.19.0); adds support for the Veracode Platform Findings REST API to access open and mitigated application findings (v26.2.18.0); introduces support for the paginated Sandbox REST API enabling retrieval of sandbox lists using pagination (v26.3.19.0); and implements a configurable retry mechanism for the Applications API (`appsec/v1/applications`) to improve handling of transient failures and timeouts (v26.2.18.0).

— Names endpoint path and wrapper versions Integration updates (2026-09-07)

03 **Healthcheck endpoint removed** BREAKING 85

The Veracode Healthcheck endpoint `/healthcheck/status` has been removed and now returns a 404 error; use the `v2/principal` endpoint in the Identity REST API instead.

— Exact old and replacement endpoints given Platform updates (2026-09-08)

04 **Flaw Importer ADO task renamed** DEPRECATED 80

The Veracode Flaw Importer task (containing spaces in the task name) is deprecated and replaced by `VeracodeFlawImporter` in the Veracode Azure DevOps extension v3.35.0; the deprecated task is scheduled for removal in a future release.

— Exact old and new task names, version given Integration updates (2026-09-07)

05 **New Reporting API fields, filters, and Audit Log report** NEW 80

Adds the `most_recent_scan` filter to the Reporting API Scans report, which when set to `true` returns only the most recent scan per asset; adds the `custom_severity` filter to the Findings report; introduces a new Audit Log report with an `action_type` filter to query user activity from Veracode Analytics; adds the `collections_list` field to the Finding Reporting API response listing collection names for an application; adds the `library_instance_ref` field to uniquely identify the library instance tied to a finding; and adds the `data_refresh_date` field indicating the most recent Analytics data refresh time.

— Names every field and filter, no endpoint path Platform updates (2026-09-08)

06 **Jira Server/Data Center 9.x support ending** DEPRECATED

75

Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends November 30, 2026 — organizations must upgrade to Jira 10 or later to maintain a supported configuration.

— Exact deadline and upgrade path given Integration updates (2026-09-07)

07 **OAuth 2.0 Client Credentials support** IMPROVED

65

Adds support for OAuth Client Credentials authentication with the Veracode XML APIs, previously limited to REST APIs, and adds the ability to create and manage OAuth 2.0 Client credentials on the API Credentials page in the Veracode Platform.

— UI path named, no endpoint or scope detail Platform updates (2026-09-08)

08 **Container and IaC scan results on the Platform** NEW

65

Makes container and IaC scan results performed from the CLI available on the Veracode Platform, with severity counts, filtering, scan history, and mitigation guidance.

— Explains what surfaces, no exact UI path Platform updates (2026-09-08)

09 **Manual Penetration Testing Flaw Matching** NEW

60

Introduces Manual Penetration Testing (MPT) Flaw Matching to correlate findings across tests for consistent flaw IDs, status values, and remediation grace period support.

— Mechanism described, no configuration steps given Platform updates (2026-09-08)

10 **Static Scan Findings filtering in VS Code** NEW

60

Adds a Current File filter to Static Scan Findings in VS Code to view findings for the file currently open in the editor, plus an option to filter and group findings by file name, with findings for each file grouped by severity (Veracode Scan for VS Code v1.16.3).

— Named UI filters, no config key given Integration updates (2026-09-07)

11 **Scan scheduling from the Platform for GitHub and Azure DevOps** NEW

60

Enables scheduling Static Analysis, SCA, and IaC scans for GitHub organizations directly from the Veracode Platform (GitHub Workflow Integration v2.0.0), and enables scheduling scans for every organization where the Azure DevOps Workflow Integration is installed, based on `veracode.yml` configuration (Azure DevOps Workflow Integration v2.3.9).

— Names scan types, no exact platform steps Integration updates (2026-09-07)

12 **GitHub app install/uninstall from the Platform** NEW

60

Enables installing and configuring the Veracode GitHub Integration directly from the Veracode Platform, eliminating manual GitHub Marketplace installation steps (GitHub Workflow Integration v2.0.0), and adds a new uninstallation workflow on the Platform to remove the Veracode app from GitHub organizations (GitHub Workflow Integration v2.1.1).

— Describes lifecycle change, no exact click path Integration updates (2026-09-07)

13 **Updated Auto-Update OWASP and CWE Top 25 rules** IMPROVED

60

Updates the Auto-Update OWASP rule and Auto-Update CWE Top 25 rule in application security policies to enforce OWASP Top 10:2025 and MITRE CWE Top 25 for 2025.

— Names exact standards and years enforced Platform updates (2026-09-08)

14 **New Analytics dashboards for IDE/CLI, Firewall, and Collections usage**

NEW

60

Introduces a new IDE and CLI Usage Explore and Dashboard in Veracode Analytics to visualize organization-wide IDE and CLI plugin usage; adds a Firewall Usage Explore and Dashboard surfacing metrics for Veracode Firewall; and adds Collections fields to the Applications explore, including new Collections visualizations in the Policy Compliance Overview dashboard.

— Named dashboards, no metrics detail Platform updates (2026-09-08)

THINNER COVERAGE BELOW

15 **Scan Scope setting in Visual Studio plugin** NEW

55

Adds a Scan Scope setting in Veracode Scan for Visual Studio to scan the full project directory or only the active solution (v0.8.0).

— Named setting with two concrete options Integration updates (2026-09-07)

16 **Usage analytics across IDE plugins** NEW

55

Adds usage analytics to Veracode Scan for Visual Studio (v0.8.0), Veracode Scan for VS Code (v1.16.0), and Veracode Scan for JetBrains (v1.8.0) to provide insights into plugin adoption, feature usage, and scan activity.

— Names three plugins and versions, no metrics shown Integration updates (2026-09-07)

17 **Disable built-in roles via Identity API** NEW

50

Adds the ability to disable Veracode-defined roles for users who use custom roles, via the Identity API.

– Names Identity API, no endpoint given `Platform updates (2026-09-08)`

18 **Autopackaging support for Apex, VB6, and Perl** NEW

50

Adds autopackaging support for Apex, VB6, and Perl in the GitHub Workflow Integration (v2.0.0).

– Names languages, no usage mechanism given `Integration updates (2026-09-07)`

19 **Region selection during VS Code OAuth authentication** NEW

45

Adds an option to change the region during OAuth authentication in VS Code, displaying the current or most recently used region (Veracode Scan for VS Code v1.16.3).

– Feature named, no further mechanism detail `Integration updates (2026-09-07)`

20 **Custom PDF REST API** NEW 45

Adds the `Custom PDF REST API` to programmatically generate custom PDF reports.

– API named, no endpoint path or parameters `Platform updates (2026-09-08)`

21 **Domain mapping to IdP for SAML SSO** NEW 40

Enables mapping multiple domains to a single identity provider (IdP) for organizations using SP-initiated SAML SSO.

– No mechanism or configuration path given `Platform updates (2026-09-08)`

22 **Visual Studio 2026 support** NEW 35

Adds Visual Studio 2026 support to Veracode Scan for Visual Studio (v0.7.2).

– Bare compatibility statement only `Integration updates (2026-09-07)`

23 **Six new built-in security policies** NEW 30

Adds six new built-in security policies to the Veracode Platform.

– No policy names or details given `Platform updates (2026-09-08)`

└──▷ ALSO FROM THESE RELEASES

. . .

└──▷ BREAKING ON UPGRADE

! The Veracode Healthcheck endpoint `/healthcheck/status` has been removed and now returns a 404 error; use the `v2/principal` endpoint in the Identity REST API instead.

! The `Veracode Flaw Importer` task (containing spaces in the task name) is deprecated and replaced by `VeracodeFlawImporter` in the Veracode Azure DevOps extension v3.35.0; the deprecated task is scheduled for removal in a future release.

! Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends November 30, 2026 — upgrade to Jira 10 or later to maintain a supported configuration.

WAS THIS USEFUL? 👍 👎

Mend.io Mend SAST

Feedback ⓘ 1 RELEASE · 2026-09-07 NOTES ↗

RANK

Mend SAST is a static application security testing service that scans source code for vulnerabilities and helps teams prioritize remediation.

> Mend SAST's biggest change this window is a unified vulnerability remediation experience with standardized fix logic that may reduce suggestion volume, alongside new multi-base-branch scanning, IDE/AI-agent suppression tooling, and JS/TS detection accuracy tuning across 11 CWEs.

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

01 Unified remediation experience with Top Fix column

BREAKING

70

Mend Platform 26.8.2 introduces a unified vulnerability remediation experience with a Top Fix column in the Findings table, an updated Remediation tab, and standardized fix logic and metadata across the Mend Platform. The updated fix logic may result in fewer remediation suggestions per vulnerability than before, so workflows relying on the previous suggestion volume will see reduced recommendations.

— Names the UI column and behavior change, but no API/config surface given. product docs

02 JS/TS detection accuracy across 11 CWEs IMPROVED

70

The JavaScript/TypeScript analysis engine's detection accuracy was improved across 11 CWEs: `CWE-22` (Path Traversal), `CWE-78` (OS Command Injection), `CWE-79` (XSS), `CWE-89` (SQL Injection), `CWE-90` (LDAP Injection), `CWE-94` (Code Injection), `CWE-117` (Log Neutralization), `CWE-295` (Certificate Validation), `CWE-601` (Open Redirect), `CWE-918` (SSRF), and `CWE-943` (NoSQL/query-language injection), including tuned handling of React controlled components to reduce false-positive DOM-based XSS findings.

— Lists all 11 CWEs precisely but offers no user action beyond scanning.

Version 26.8.2 (07-September-2026)

03 Suppression tool in IDE and AI agent workflow NEW

60

A new suppression tool in the IDE and AI agent workflow supports single and bulk suppression or unsuppression of findings across all scanning engines, including SAST and SCA, with UI updates reflecting within seconds.

— Describes scope and speed but no exact command or menu path. product docs

04 **Multi-base-branch scanning on Azure DevOps and Bitbucket Cloud**

NEW

60

The Developer Platform now supports scanning multiple base branches into dedicated projects on Azure DevOps and Bitbucket Cloud, enabling security managers to monitor and set policies for each branch individually.

— Names platforms and effect but no configuration details given. product docs

└──▷ **BREAKING ON UPGRADE**

! The updated fix logic may result in fewer remediation suggestions per vulnerability; existing workflows relying on the previous volume of remediation recommendations will see reduced suggestions.

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ 1 RELEASE · 2026-09-07 BLOG ↗ RANK

StepSecurity hardens CI/CD pipelines by detecting and preventing supply chain attacks through runtime security monitoring and vulnerability detection.

> StepSecurity's only shipped change this window expands its Webhook Integration from a single tenant-wide webhook to up to ten independently configured endpoints, each with its own event subscriptions and repository scope.

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

01 **Multiple webhook endpoints per tenant** IMPROVED 82

StepSecurity now supports up to ten webhooks per tenant, each with its own endpoint URL, HTTP method, headers, message format, and `Envelope` template — replacing the previous single-webhook-per-tenant limit. Each webhook can subscribe to any combination of four event types (`Insights` , `Detections` , `Secure Registry Audit Logs` , `API Audit Logs`) and can be scoped to a subset of repositories or to all current and future repositories, with independent delivery per endpoint so a failure at one does not affect others. The UI adds an `Add Webhook` control that creates additional endpoints, each shown as a collapsible card with name, endpoint, and event-type chips.

Route runtime detections to your SIEM and audit logs to a separate log archive without either receiver filtering unwanted traffic.

> 📍 In the StepSecurity console, go to Webhook Integration and click `Add Webhook` . Configure the first webhook with endpoint URL pointing to your SIEM, subscribe it to `Detections` and `Insights` , and set the Envelope template and auth headers your SIEM expects. Click `Add Webhook` again, set the second endpoint URL to your log archive, subscribe it to `Secure Registry Audit Logs` and `API Audit Logs` , then click `Save changes` .

— Names event types and UI control but no API endpoint given

`Webhook Integration: Multiple Endpoints…`

WAS THIS USEFUL? 👍 👎

Mend.io Mend Platform

Feedback ⓘ 1 RELEASE · 2026-09-06

NOTES ↗ RANK

Traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend shipped a unified vulnerability remediation experience with a new Top Fix column, multi-base-branch SAST scanning for Azure DevOps and Bitbucket Cloud, an in-IDE/AI-agent suppression tool, JavaScript/TypeScript detection accuracy tuning across 11 CWEs, and several Renovate manager and datasource additions.

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

01 JS/TS detection accuracy improvements across 11 CWEs IMPROVED

70

Improves detection accuracy of the JavaScript/TypeScript analysis engine across CWE-22, CWE-78, CWE-79, CWE-89, CWE-90, CWE-94, CWE-117, CWE-295, CWE-601, CWE-918 and CWE-943, including tuning for React controlled components to reduce false positives where query parameters flowing into escaped JSX attributes were incorrectly flagged as DOM-based XSS.

— Detailed mechanism and named CWEs but nothing for a reader to invoke product docs

02 Renovate manager and datasource additions NEW

70

Renovate now passes `--no-frozen-lockfile` to pnpm when performing lock file updates in the `manager/npm` manager, supports optional channel options in the `manager/pixi` manager, and supports AWS S3 as a Helm charts repository in the `datasource/helm` datasource.

— Names exact flags, managers and datasource for direct config use product docs

03 Unified vulnerability remediation experience BREAKING

67

Adds a Top Fix column to the Findings table and an updated Remediation tab, standardizing fix logic and recommendation metadata across SAST and SCA findings on the Mend Platform. Because the new fix logic is more accurate, it may surface fewer remediation suggestions per vulnerability than before, reducing noise but potentially removing previously shown suggestions — existing workflows relying on the old suggestion volume should be reviewed.

— Names UI surfaces and behavior change but no config/API to act on

Version 26.8.2 (06-September-2026)

04 **Suppression tool for IDE and AI agent workflows** NEW

63

Adds a suppression tool in Agentic Integrations supporting single and bulk suppression or unsuppression actions across all scanning engines, including SAST and SCA, with UI updates reflecting within seconds to reduce context-switching for developers working in their IDE or AI agent workflow.

— Explains behavior and scope but only a UI entry point, no exact command

Version 26.8.2 (06-September-2026)

---

05 **Multi-base-branch scanning in Mend SAST** NEW

60

Mend SAST now supports scanning multiple base branches into dedicated projects in the Developer Platform on Azure DevOps and Bitbucket Cloud, letting security managers monitor and set policies for each branch individually.

— Clear capability and scope but only a UI/product-area starting point product docs

└──▷ **BREAKING ON UPGRADE**

! The unified fix logic may produce fewer remediation suggestions per vulnerability due to increased accuracy, reducing overall noise but potentially removing previously shown suggestions.

! The unified vulnerability remediation fix logic may produce fewer remediation suggestions per vulnerability; existing workflows that relied on the previous volume of suggestions should be reviewed.

WAS THIS USEFUL? 👍 👎

ZeroPath

Feedback ⓘ 1 RELEASE · SEEN 2026-09-08 NOTES ↗ RANK

AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives. Automated vulnerability detection, instant fixes, and seamless CI/CD integration for modern DevOps teams.

> ZeroPath consolidated all scanning configuration into a new Scanner workspace with scope-aware permission enforcement, added SLO/SLI tracking and an audit trail for service levels, and introduced an org-wide toggle to redact AI-generated descriptions for secret findings.

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

01 Scope-aware scanner permissions and authorization checks

88

BREAKING

Adds scoped permissions controlling `Rules`, repository context, and scan locations that surface only the actions a member or team is authorized to perform, with inline reasons when `Add`, `Save`, or `Delete` are unavailable. Editing or deleting a rule, repository context, or scan location set is now authorized against both the scope a definition already holds and the scope being assigned, blocking moves outside a member's granted repositories. The teams permission editor lists every grantable permission, and an organization-wide repository permission grant carries its named actions — including saving repository context or scan location sets — across all repositories; members with per-repository scanner permission see a searchable repository list under `Scanner` scoped to repos they can configure, and definitions linked to a repository removed from the version control host remain deletable.

Grant a team organization-wide repository permissions so automation can save scan location sets across all repositories without per-repo configuration.

> 📍 In the permission editor, open the team's grant, select the repository permission that includes 'save scan location set', and enable the organization-wide grant — automation using that team's credentials will then be authorized to write scan location sets across all repositories.

Let a developer with per-repository scanner access configure only their own repos without exposing organization-wide settings.

> 📍 In the main sidebar, click `Scanner` — members with per-repository scanner permission will see a searchable list scoped to only the repositories they hold permission on, with Add/Save/Delete showing inline reasons when unavailable.

— Details exact permission mechanism, actions, and a breaking authorization change

Scanner workspace, grouped settings nav…

02 Unified Scanner workspace with grouped settings navigation

NEW

78

Adds a dedicated Scanner workspace in the main sidebar that consolidates all scanning configuration under Configuration (scanner settings, scan locations) and Detection (repository context, rules, sources and sinks). Settings sidebar pages are now grouped by purpose — Workspace (general, members, teams, billing), Security (scanner settings, repository tags, service levels), and Connections (integrations, notifications, API keys) — and scanner settings appear across five consistent categories, General, Full scans, Pull requests, Dependencies, and Remediation, at organization, tag, and repository scope so overrides appear in the same place at every level.

Navigate to scanner detection rules for a specific repository using the new Scanner workspace sidebar.

> 📍 In the main sidebar, select Scanner › Detection › Rules, then pick the target repository or tag scope from the scope selector to view or edit rules scoped to that repository.

Override pull request scan settings at the repository level without losing your place in the settings hierarchy.

> 📍 In Settings › Security › Scanner settings, select the target repository scope, open the Pull requests category, adjust the threshold, and switch to another category — the edit is preserved until you save.

— Names workspace areas and categories, gives exact navigation paths

Scanner workspace, grouped settings nav…

03 SLO/SLI tracking and audit trail for service levels NEW

70

Adds Settings → Service levels → Policies to configure SLO targets on individual policies, a `Settings → Service levels → SLO / SLI` section with trend charts turning SLA activity into a time-series view, and an Audit trail section under Settings → Service levels to record SLA-related activity.

Set an SLO target on a policy and monitor compliance trends over time.

> 📍 In the console, go to Settings → Service levels → Policies, select a policy, and set an SLO target. Then navigate to Settings → Service levels → SLO / SLI to view trend charts tracking compliance against that target.

— Names exact settings paths and sections, no deeper mechanism given product docs

04 Org-wide toggle to redact secret finding descriptions NEW

60

Adds a Hide Descriptions toggle in the Repositories & data section for organization admins, which redacts AI-generated descriptions and exploit walkthroughs for secret findings across the dashboard, API responses, webhook payloads, notifications, and exported tickets.

Prevent prose that might quote a secret value from appearing in dashboards, API responses, webhooks, and exported tickets — useful when sharing findings with contractors or exporting to third-party systems.

> 📍 In the ZeroPath console, go to Repositories & data and click Hide Descriptions to enable the org-wide redaction toggle.

— Names the toggle and every affected surface, gives UI path only `product docs`

└──▷ **BREAKING ON UPGRADE**

! Editing or deleting a rule, repository context, or scan location set is now authorized against the scope the definition already holds as well as the scope being assigned — members and automation moving a definition outside their granted repositories will be refused.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗ RANK

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

> Elcomsoft Quick Triage 2.2 introduces a merged timeline view, a new plugin-based architecture for artifact support, and a file system snapshot artifact, alongside extended password recovery, full-text search, and BitLocker key extraction capabilities.

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

01 **Plugin architecture for artifact support** IMPROVED 40

The tool introduces a plugin architecture that rebuilds artifact support as a plugin system, replacing the prior monolithic approach.

— Names the change but not plugin API or usage steps. changelog-20260908-262ddce6

02 **Unified timeline view** NEW 35

Quick Triage 2.2 adds a timeline view that merges events from different artifacts into a single chronology.

— Describes function but no mechanism or UI detail. changelog-20260908-262ddce6

03 **Extended password recovery, search, and BitLocker key extraction**

IMPROVED 25

This release extends password recovery capabilities, extends full-text search capabilities, and extends BitLocker key extraction capabilities.

— Three named areas extended but no specifics given. changelog-20260908-262ddce6

04 **File system snapshot artifact** NEW 20

Adds file system snapshot as a new artifact type supported by Quick Triage.

— Bare mention with no detail on contents or use. changelog-20260908-262ddce6

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Snyk Container

Feedback ⓘ 1 RELEASE · 2026-09-07 BLOG ↗

RANK

Snyk Container scans Docker images and container registries for vulnerabilities, misconfigurations, and malware to secure containerized applications.

> Snyk Container now surfaces source repository information for locally built images by reading BuildKit metadata.

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

01 **Source repository info via BuildKit metadata** NEW 38

Container scans now surface source repository information for locally built images using BuildKit metadata.

— Names mechanism (BuildKit metadata) but no command or config shown. 2026-09-07-1c4c9dfc

WAS THIS USEFUL? 👍 👎

Inspektor Gadget

Feedback ⓘ 1 RELEASE · 2026-09-07 NOTES ↗

CODE ↗ RANK

Inspektor Gadget is a set of tools and framework for data collection and system inspection on Kubernetes clusters and Linux hosts using eBPF

> Inspektor Gadget v0.56.0 introduces experimental Kubernetes multi-tenancy with RBAC enforcement, moves OCI image management to a rootless per-user store, and extends uprobe tracing to stripped Go binaries.

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

01 **Rootless OCI image store and commands** BREAKING 83

Adds a `--oci-store-user` flag to `sudo ig run` so gadgets can be read from a per-user OCI store at `~/.ig/oci-store` (built rootlessly) without creating root-owned files in the user home directory, and enables `ig image build`, `ig image list`, `ig image inspect`, `ig image push`, and `ig image pull` to run as a normal non-root user against that same store. This is a breaking change: gadgets previously built into the root-owned store will not be visible to rootless `ig image` commands.

— Names exact flag, path, and all five commands with migration note v0.56.0

02 **Uprobe support for stripped Go binaries** NEW 65

The `uprobetracer` now falls back to Go's `.gopclntab` section (including the Go 1.26+ `.go.module` format) when the ELF symbol table is absent, enabling uprobes on stripped Go binaries such as a production kubelet.

— Explains mechanism and gives concrete example, no runnable command v0.56.0

03 **Kubelet socket auto-detection via /configz endpoint** BREAKING 65

Container runtime socket auto-detection now uses the kubelet's `/configz` endpoint over HTTPS, requiring only the `nodes/configz` RBAC permission instead of the broader `nodes/proxy` (Kubernetes v1.33+; earlier clusters fall back to the configured socket path). This is a breaking change: v1.33+ deployments must update RBAC rules to grant `nodes/configz` instead of `nodes/proxy`.

— Names exact RBAC permissions and version requirement, clear migration v0.56.0

THINNER COVERAGE BELOW

04 **Kubernetes multi-tenancy with RBAC enforcement** NEW 57

An experimental multi-tenancy mode authenticates every `kubectl-gadget` request via `TokenReview` and restricts gadgets to RBAC-authorized namespaces using `SubjectAccessReview`, enforced server-side in the gadget service via KubeManager.

— Names mechanism (TokenReview/SubjectAccessReview) but no setup steps v0.56.0

---

05 **Network namespace in trace_mount output** IMPROVED 25

The `trace_mount` gadget now reports network namespace in its output.

— One-line addition with no further detail given v0.56.0

└──▷ **BREAKING ON UPGRADE**

! Kubernetes deployments on v1.33+ now require the `nodes/configz` RBAC permission instead of `nodes/proxy`; clusters that granted only `nodes/proxy` must update their RBAC rules.

! `ig image build`, `list`, `inspect`, `push`, and `pull` now use a per-user OCI store at `~/.ig/oci-store` instead of a root-owned store; gadgets previously built into the root store will not be visible to rootless `ig image` commands.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ 1 RELEASE · 2026-09-08 NOTES ↗ RANK

Plerion is a cloud security platform that discovers and manages vulnerabilities and misconfigurations across infrastructure.

> Plerion's Kubernetes collector manager is now a multi-architecture image, adding native arm64 (Graviton) support alongside amd64.

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

01 **Multi-arch Kubernetes collector manager image** IMPROVED

62

The Kubernetes collector manager is now published as a multi-architecture image supporting both `amd64` and `arm64`, enabling deployment on Graviton node pools without requiring a dedicated x86 node or `nodeSelector` configuration.

— Names exact architectures but no rollout/version specifics

The Kubernetes collector manager runs o…

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

Tenzir

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗ RANK

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

> Tenzir added Splunk HEC acknowledgement support to the Tenzir Node, enabling confirmed event delivery.

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

01 **Splunk HEC acknowledgement support** NEW 38

The Tenzir Node now supports Splunk HTTP Event Collector (HEC) acknowledgements, enabling confirmed event delivery over the HEC protocol.

— Names the protocol and mechanism but no config detail `changelog-20260908-4b6c0ef0`

WAS THIS USEFUL? 👍 👎

◆ SIEM

Google SecOps

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗

RANK

Meet your business challenges head on with AI and cloud computing services from Google, including security, data management, and hybrid & multi-cloud.

> Google SecOps added self-service license download for Bindplane Enterprise (Google Edition) and introduced v2 feed types built on Cloud Storage Transfer Service, with a planned migration deadline for legacy feeds.

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

01 Self-service Bindplane Enterprise license download NEW

65

Enterprise Plus and Google Unified Security (GUS) customers can now download the Bindplane Enterprise (Google Edition) license key directly from the platform console, under SIEM Settings > Collection Agents > Bindplane Enterprise (Google Edition), to activate a self-managed deployment.

Download your Bindplane Enterprise license key from the Google SecOps console to activate a self-managed Bindplane Enterprise (Google Edition) deployment.

> 📍 In the Google SecOps console, navigate to SIEM Settings › Collection Agents › Bindplane Enterprise (Google Edition) and download your license key.

— Names exact console path and purpose, no deeper mechanism given

changelog-20260908-ef1f798a

02 v2 feed types via Cloud Storage Transfer Service NEW

62

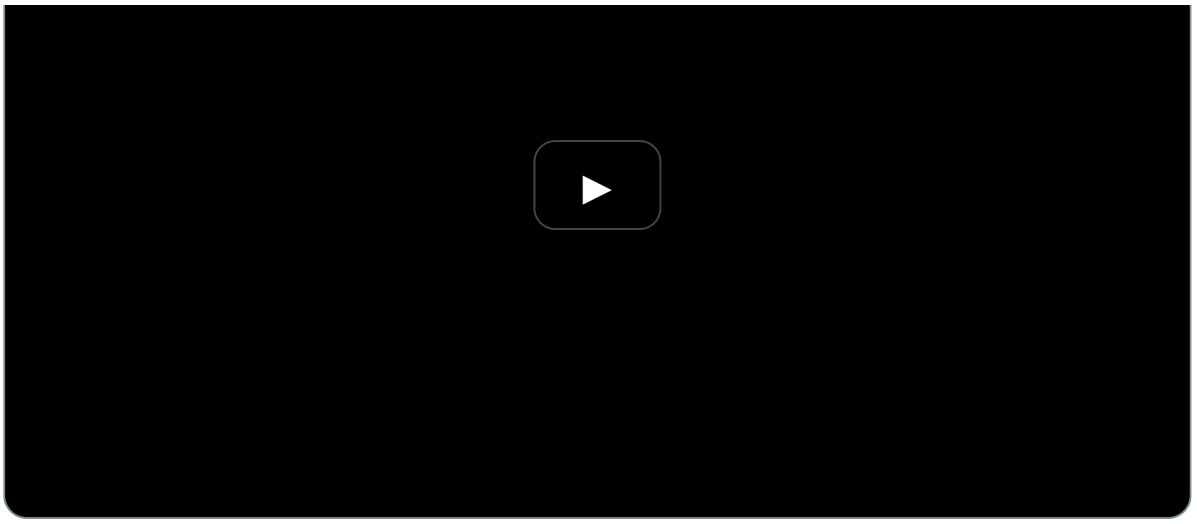
New v2 feed types use the Google Cloud Storage Transfer Service (STS) for improved performance, scalability, and reliability. Feeds using v1 feed types will be automatically migrated to v2 starting April 6, 2026, and some feeds may require additional IP allowlist or service account permission updates before that migration.

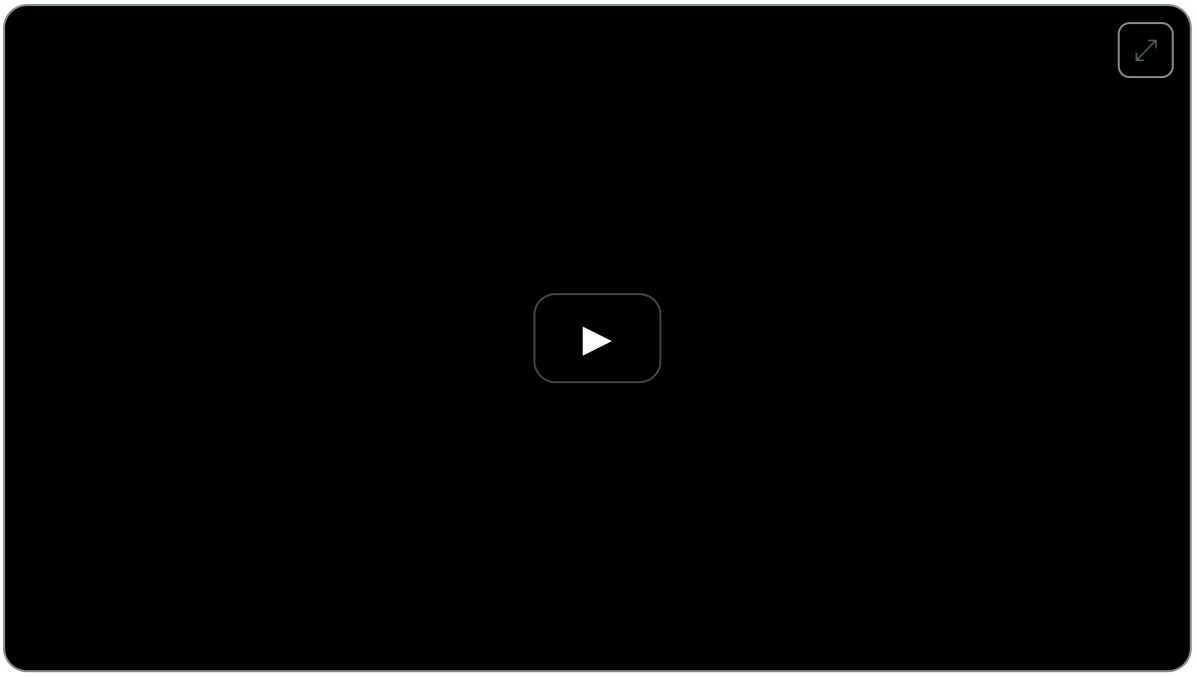
— Names mechanism and migration date but no config steps given

changelog-20260908-ef1f798a

└──▷ **ALSO FROM THESE RELEASES**







└──▷ **BREAKING ON UPGRADE**

! Feeds using v1 feed types will be automatically migrated to v2 starting April 6, 2026; some feeds may require additional IP allowlist or service account permission updates before migration.

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗ RANK

Databricks is a unified analytics platform for data engineering, data science, and machine learning built on Apache Spark.

> Databricks brings automatic change data feed to general availability, letting teams track row-level changes at query time without per-table setup and with faster merge/update performance.

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

01 **Automatic change data feed (Auto CDF) GA** NEW 83

Auto CDF computes row-level changes at query time using row tracking, eliminating the need to manually run `ENABLE CHANGE DATA FEED` on each Delta Lake or Apache Iceberg v3 table. It makes `MERGE` and `UPDATE` operations approximately 15% faster on tables queried for changes by removing write-time overhead, and works with batch queries, Structured Streaming, and Delta Lake Sharing when using Databricks Runtime 19 LTS or above with row tracking enabled.

— Names mechanism, perf gain, and runtime requirement but no exact setup command

Automatic change data feed is now gener…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗ RANK

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

> DNSFilter overhauled its dashboard experience this window, making Dashboards the default landing view, merging its filter systems, adding live status indicators to deployment widgets, and clarifying TLD labeling in lists and logs.

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

01 **Unified Filters experience in dashboard** IMPROVED 65

Quick Filters and Advanced Filters are merged into a single Filters experience in the dashboard Filters drawer, grouped by Source (Organizations, Sites, Roaming Clients/Relays, Users) and Traffic (Result).

– Names exact grouping categories and UI location changelog-20260908-ec47516d

02 **Live status indicators on deployment widgets** NEW 60

Dashboard deployment widgets (Roaming Clients Protected, Sites Protected, Relays, Users, Collections) now show a pulsating green dot with a tooltip indicating the data is live and not affected by the time range filter.

– Names all affected widgets and exact indicator behavior changelog-20260908-ec47516d

THINNER COVERAGE BELOW

03 **Top Level Domain labeling for bare TLD entries** IMPROVED 55

Bare TLD entries (e.g. `.ru`, `.xyz`) in Allow/Block lists and the Query Log now display as Top Level Domain instead of Uncategorized, distinguishing intentional TLD-level rules from genuinely uncategorized domains.

– Names the label change and where it appears, with examples changelog-20260908-ec47516d

04 **Dashboards set as default landing view** BREAKING 50

Dashboards is now the default landing view for all users; the legacy Overview page is retired and existing bookmarks redirect automatically to Dashboards.

– Names the change and effect but no migration detail beyond redirect

changelog-20260908-ec47516d

WAS THIS USEFUL?  

Firezone

Feedback ⓘ 3 RELEASES · 2026-09-07 → 2026-09-08 CHANGELOG ↗

RANK

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

> Firezone added X.509 certificate-backed device trust across its Clients, retired the `FIREZONE_NAME` / `-n` Gateway option in favor of naming Gateways in the Admin Portal, and made smaller improvements to Gateway flow-log storage, device reporting, and update notifications.

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

01 Gateway naming moved to Admin Portal BREAKING

85

The `FIREZONE_NAME` environment variable and `-n` CLI flag have been removed from the Gateway; the Gateway's name must now be configured in the Admin Portal. This is a breaking change for any deployment scripts or container configs that set `FIREZONE_NAME` or pass `-n`.

– Names exact flag/env var and required migration step. changelog-20260908-1753f87c

Gateway 1.6.1 (2026-09-07)

THINNER COVERAGE BELOW

02 Gateway flow-log storage and device reporting IMPROVED

55

Gateway 1.6.1 auto-creates the `/var/lib/firezone` directory to ensure flow logs can be persisted, and now reports the device serial and UUID.

– Names the directory and new reported fields, no user action needed.

Gateway 1.6.1 (2026-09-07)

03 Client update notification behavior improved IMPROVED

45

Update notifications now appear once per Client launch instead of only once per new version; on Linux, update notifications also trigger the download of the new package immediately upon appearing.

– Describes behavior change but no config or command given. changelog-20260908-1753f87c

Linux GUI 1.5.17 (2026-09-07) Windows GUI 1.5.17 (2026-09-07)

04 X.509 certificate-backed device trust NEW

30

Adds support for device trust backed by X.509 certificates, shipped in the Linux GUI, Windows GUI, and Windows Headless clients.

— No mechanism, enrollment, or config details given. `changelog-20260908-1753f87c`

`Linux GUI 1.5.17 (2026-09-07)` `Windows Headless 1.5.12 (2026-09-07)`

└─▷ BREAKING ON UPGRADE

! The `FIREZONE_NAME` flag (`-n`) has been removed from the Gateway; Gateways that set their name via this option must now set the name in the Admin Portal instead.

! The `FIREZONE_NAME` environment variable and `-n` CLI flag have been removed from the Gateway; any deployment scripts or container configs that set `FIREZONE_NAME` or pass `-n` will break — the Gateway name must now be configured in the Admin Portal.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Auth0

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗ RANK

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

> Auth0's biggest update lets Organization Admins self-serve Cross-App Access configuration via API, brings Advanced MFA Configurations to General Availability, and opens Early Access for apps to request third-party tokens and for a redesigned dashboard navigation.

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

01 Advanced MFA Configurations reaches General Availability

IMPROVED 95

Advanced MFA Configurations is now GA, enabling configurable Remember this device idle lifetimes (1 hour to 30 days) and absolute lifetimes (up to 90 days), OTP lengths for SMS and email factors (4–10 digits), OTP lifetime for SMS MFA (30 seconds to 1 hour), and OTP lifetime for email MFA (5 minutes to 1 hour). Configured via Dashboard under Security > Multi-factor Auth by setting OTP Length and OTP Lifetime per factor.

Tighten MFA UX by reducing OTP length to 4 digits and capping SMS OTP lifetime to 2 minutes for a lower-friction, shorter-window factor.

> 📍 In the Auth0 Dashboard, go to Security › Multi-factor Auth, open the SMS/Voice or Email factor settings, set OTP Length to 4, and set OTP Lifetime to 120 seconds.

— Concrete numeric ranges plus exact dashboard steps. changelog-20260908-34570346

02 Cross-App Access Requesting App (Early Access) NEW

81

Auth0 now supports acting as a requesting application in Cross-App Access flows (Early Access), built on Auth0 Token Vault using the ID-JAG (Identity Assertion Authorization Grant) protocol, letting apps fetch third-party API tokens for enterprise users without a manual Connect your account step. Supported on OIDC and Okta Workforce connections with org-aware resolution via `org_id`; a new Dashboard toggle enables Cross-App Access Requesting App on enterprise connections.

Enable Cross-App Access Requesting App on an enterprise OIDC connection so your app can fetch third-party API tokens for org users without a manual consent step.

> 📍 In the Auth0 Dashboard, go to Authentication › Enterprise Connections, select your OIDC or Okta Workforce connection, and toggle Cross-App Access Requesting

App to enabled.

— Names protocol and connection types; toggle path is UI-only.

changelog-20260908-34570346

03 **Self-service XAA configuration via My Organization API** NEW

65

Adds self-service configuration of Third-Party Application Access and Cross App Access (XAA) through the My Organization API and Embeddable UI Components (Open Early Access), letting Organization Admins configure third-party app access, connection-level third-party access, and XAA Resource Application roles without routing changes through the Auth0 Dashboard.

— Names the API and scope but no exact endpoint path. changelog-20260908-34570346

THINNER COVERAGE BELOW

04 **Dashboard Navigation and IA Refresh (Early Access)** NEW

45

A redesigned sidebar, improved information architecture, consistent breadcrumbs, and a modernized visual experience, toggled per user account from the profile menu.

— Describes visual changes but only a vague toggle location. changelog-20260908-34570346

WAS THIS USEFUL?

Ping Identity

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗

RANK

Read docs, explore use cases, learn best practices

> Ping Identity's rapid channel release adds IGA identity and role composition certification with dedicated review views, correlation rules for unmanaged applications, Friendly Captcha and no-JS QR code options for MFA registration, alongside breaking changes to relationship validation and Google Secret Manager secret identifiers.

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

01 Unique JWK kid identifiers for Google Secret Manager

BREAKING

90

ESV secrets are now stored in Google Secret Manager using a unique name-and-version identifier as the default JWK `kid`; deployments where multiple secrets shared the same GSM version must set the new `esv.am.secrets.gsm.stableid.version.only` environment variable (`esv.am.secrets.gsm.stableid.version.only=true`) to retain the legacy `kid` value.

Preserve the legacy GSM JWK `kid` format in environments where existing clients depend on the old version-only identifier.

```
$ esv.am.secrets.gsm.stableid.version.only=true
```

— Names exact env var, value, and mitigation path. changelog-20260908-74d79c95

02 Role composition certification in IGA NEW

62

Adds IGA role composition certification, enabling reviewers to certify role definitions including attributes, glossary data, entitlements, applications, and membership rules, with a dedicated review task view for role-composition items.

— Lists concrete role elements covered by certification. changelog-20260908-74d79c95

THINNER COVERAGE BELOW

03 Stricter relationship property validation

BREAKING

58

Advanced Identity Cloud now rejects relationship creation requests that use a property that does not accept the referenced object type, instead of silently creating the relationship through a different valid property; existing integrations relying on the silent fallback behavior will now receive errors.

— Explains before/after but no endpoint or field named. changelog-20260908-74d79c95

04 **No-JS QR codes for MFA registration nodes** NEW 57

Adds an option to the Combined MFA Registration, OATH Registration, and Push Registration nodes to render QR codes without requiring JavaScript.

— Names three specific nodes gaining the option. changelog-20260908-74d79c95

05 **Identity profile certification in IGA** NEW 55

Adds IGA identity profile certification, enabling reviewers to certify identity records including external users and AI agents — not just access — with a dedicated review task view for identity-profile certification items.

— Describes scope but no UI path or config named. changelog-20260908-74d79c95

06 **Friendly Captcha in v2 CAPTCHA node** NEW 53

Adds support for Friendly Captcha to the `v2 CAPTCHA node` as a privacy-focused CAPTCHA provider alternative.

— Names the node but no mechanism or config detail given. changelog-20260908-74d79c95

07 **Role request membership handling in IGA** IMPROVED 50

IGA role request forms now display current role membership as read-only, and each member addition or removal in the role lifecycle management screen is submitted as an individual grant or removal request.

— Explains before/after behavior but not exact screen paths. changelog-20260908-74d79c95

08 **Correlation rules for unmanaged apps in IGA** NEW 45

Adds IGA correlation rules for unmanaged applications, letting administrators link aggregated accounts to identities during account aggregation.

— Describes purpose but not rule syntax or UI location. changelog-20260908-74d79c95

09 **IGA activity and notification monitoring GA** NEW 45

Adds IGA activity and notification monitoring as generally available, with event type filtering and audit log presentation.

— States GA status and filtering but no config surface. changelog-20260908-74d79c95

10 **Additional IGA administration capabilities** NEW 40

IGA administration gains the ability to delete certificate templates via the IGA UI, to change a campaign's title and owner after the campaign has been launched, and to create approval tasks automatically when expiration is configured with a workflow variable.

— Three thin admin additions with no further mechanism. changelog-20260908-74d79c95

11 **Consistent justification enforcement in IGA** IMPROVED

40

IGA now enforces the required request justification setting consistently across lifecycle management, request extension, and entitlement-grant requests.

— Names affected request types but no config key. changelog-20260908-74d79c95

12 **Managed user updates during attribute conflicts** IMPROVED

40

Enables updating a managed user even when the change conflicts with generated attribute values, propagating the update to all associated workforce application accounts during queued synchronization.

— Describes behavior change but no named flag or endpoint. changelog-20260908-74d79c95

└──▷ **BREAKING ON UPGRADE**

! Advanced Identity Cloud now rejects relationship creation requests that use a property that does not accept the referenced object type, instead of silently creating the relationship through a different valid property — existing integrations relying on the silent fallback behavior will now receive errors.

! ESV secrets are now stored in Google Secret Manager using a unique name-and-version identifier as the default JWK `kid`; deployments where multiple secrets shared the same GSM version must set `esv.am.secrets.gsm.stableid.version.only` to retain the legacy `kid` value.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Certiv

Feedback ⓘ 1 RELEASE · 2026-09-04 NOTES ↗ RANK

Complete visibility and control over every AI agent locally installed on the endpoint. Discover, understand, and safeguard every AI action across your organization.

> Certiv launched Certiv Cost in public preview, adding session-level token attribution, per-endpoint dollar budgets with tiered enforcement and rate calibration, and end-user spend visibility in Scout.

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

01 Per-endpoint budget enforcement modes and rate calibration

NEW 95

Certiv Cost adds per-endpoint dollar budgets with three enforcement modes: Report (logs overage, no interruption), Notify (fires alerts, work continues), and Block (stops future prompts at the endpoint before the request leaves). Organizations can enter negotiated per-model rates so thresholds are computed against actual cost rather than list pricing, and a budget simulation feature runs a proposed threshold against historical usage to project how often it would have been triggered before enforcement is turned on.

Start in Notify mode to understand real consumption patterns across your endpoints before committing to hard limits that could block legitimate work.

> 📍 In the Certiv console, go to Cost › Endpoints, select an endpoint, set a dollar budget for your desired period, and choose Notify as the enforcement mode — alerts will fire when the threshold is crossed but prompts will continue to go out.

Switch to Block mode on high-risk endpoints once you have a baseline, so a runaway agent session is stopped at the endpoint before the prompt leaves and drains shared quota.

> 📍 In the Certiv console, go to Cost › Endpoints, select the endpoint, and change the enforcement mode from Notify to Block — future prompts exceeding the budget are stopped before the request is sent to the model.

Use the budget simulation against historical usage to pick a threshold that would have caught past overruns without generating noise before switching an endpoint to Block mode.

> 📍 In the Certiv console, go to Cost › Endpoints, select an endpoint, enter a proposed dollar budget, and run the simulation — Certiv Cost shows how many

times that ceiling would have been triggered over the historical period.

— Names exact modes and gives runnable console navigation steps.

Token Spend Is a Potential Signal for a…

02 **Session-level token attribution in Certiv Cost** NEW

60

Certiv Cost enters public preview with session-level token attribution that tracks spend by agent, endpoint, user, and model, plus historical trend views to surface abnormal consumption.

— Explains dimensions tracked but no concrete steps given.

Token Spend Is a Potential Signal for a…

THINNER COVERAGE BELOW

03 **End-user spend visibility in Scout** NEW

40

Scout now shows each user their own token consumption on their own endpoint, broken out by hour.

— One-line description with no navigation or mechanism detail.

Token Spend Is a Potential Signal for a…

WAS THIS USEFUL? 👍 👎

ToolHive

Feedback ⓘ 1 RELEASE · 2026-09-08 NOTES ↗ RANK

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive v0.47.0 strengthens its trust and identity model with SPIFFE-based operator CRDs and cosign key-pinned verification for skill and plugin installs, alongside a batch of smaller auth-hardening and proxy tweaks.

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

01 **Key-signed verification for skill and plugin installs** NEW

75

Skill and plugin installs are now verified against a public key using cosign, with the pinned cosign key recorded in the lock schema. The pinned key is honoured on `sync` and `upgrade` operations, and plugin pushes are signed with the lock feature gate removed.

— Names lock schema and `sync`/`upgrade` commands explicitly. v0.47.0

THINNER COVERAGE BELOW

02 **SPIFFE trust configuration in operator CRDs** NEW 45

Adds SPIFFE trust configuration and exposes SPIFFE client-auth registration through the operator CRDs, allowing workload identity to be managed declaratively.

— Names SPIFFE and CRDs but no field or command shown. v0.47.0

03 **Private CA trust for RFC 8693 trusted issuers** NEW

45

Trusts private CAs for RFC 8693 trusted issuers.

— Names the RFC and issuer mechanism but no config key. v0.47.0

04 **OIDC discovery guarded against private IPs** IMPROVED

35

Guards OIDC discovery against private IPs.

— Names OIDC discovery but no config or example. v0.47.0

05 **Cedar delegation without upstream session** IMPROVED

30

Allows Cedar policy delegation to occur without requiring an upstream session.

— States the capability but no mechanism or config detail. v0.47.0

06 **Canonical inbound grants exposed in CRDs** NEW 30

Exposes canonical inbound grants in the operator CRDs.

— Names CRD surface but no field or example. v0.47.0

07 **Claude workflows routed through Stacklok LLM gateway** NEW 30

Routes Claude workflows through the Stacklok LLM gateway.

— Names the gateway but no setup or routing detail. v0.47.0

08 **OAuth callback Referer header suppressed** IMPROVED 30

Suppresses the OAuth callback Referer header.

— Names the exact header but no broader mechanism given. v0.47.0

09 **Configurable proxy request-body limit** IMPROVED 25

Makes the proxy request-body size limit configurable.

— No flag or config key given for the new limit. v0.47.0

10 **Embedded auth servers without upstream MCP servers** IMPROVED 25

Allows embedded auth servers to operate without upstreams.

— Bare statement with no mechanism or config detail. v0.47.0

11 **GUI editor detection without PATH binary** IMPROVED 25

Detects GUI editors even when no corresponding binary is present on the PATH.

— Bare behavioural note with no further detail. v0.47.0

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 1 RELEASE · 2026-09-08 CHANGELOG ↗ RANK

Connect, protect, and build everywhere.

> Cloudflare's biggest changes this window are a 64 MiB unified Worker bundle limit with new deploy-time inspection flags and WSGI/ASGI framework support in Python Workers, alongside a substantial expansion of the Images binding and several Cloudflare One Appliance dashboard and networking improvements.

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

01 Bundle size inspection and unified 64 MiB Worker limit

IMPROVED 90

Adds `--outdir` and `--dry-run` flags to `wrangler deploy` to inspect uncompressed bundle size before deploying, now checked against a unified 64 MiB limit across Free and Paid plans (previously 3 MB / 10 MB compressed).

Check a Worker's uncompressed bundle size against the 64 MiB limit before deploying to catch oversized builds in CI without actually publishing.

```
$ wrangler deploy --outdir bundled/ --dry-run
```

— Exact flags, before/after limits, and a runnable command. changelog-20260908-46bc8c4b

02 WSGI/ASGI framework support in Python Workers NEW

80

Adds wsgi.entrypoint() and asgi.entrypoint() from the `workers` module to support WSGI frameworks (Django, Flask) and ASGI frameworks (FastAPI, Starlette) in Python Workers.

Run a Django app inside a Python Workers by wiring Django's WSGI application into the Workers runtime entry point.

```python
from workers import wsgi
from django.core.wsgi import get_wsgi_application

app = get_wsgi_application()
Default = wsgi.entrypoint(app)
```

— Named entrypoints and frameworks with a working code sample.

changelog-20260908-46bc8c4b

03 **Images binding gains text, metadata filtering, signed URLs and uploads**

NEW 80

The Images binding adds .text() to rasterize text into or over an image with `font`, `size`, and `color` style options (plus a `text` key in the `draw` array of `cf.image`); `filter.metadata` on .list() to filter hosted images by custom metadata including bounded range conditions like `priority: { gte: 2, lte: 5 }`; .signedUrl() to generate signed URLs for private hosted images server-side; .createDirectUpload() to create Direct Creator Upload links for client-side uploads; and a `headers` option on .response() to set response headers in one call, noting Content-Type is always taken from the optimized image and cannot be overridden.

— Every new method and option named verbatim, no usage example given.

changelog-20260908-46bc8c4b

04 **Custom application traffic steering in Appliance dashboard**

NEW 75

Adds a Source subnets field in Traffic Steering (Breakout traffic / Prioritized traffic) on the Cloudflare One Appliance dashboard to match traffic by source IP range when defining custom applications, and enables defining and editing custom applications directly from the dashboard under Traffic Steering without calling the API.

Add a source-subnet match to a custom breakout application on Cloudflare One Appliance so that traffic from a specific internal IP range is steered without an API call.

> 📍 In the Cloudflare One Appliance dashboard, go to Traffic Steering › Breakout traffic, select Assign application traffic › Add, enter hostnames or IP subnets, then set Source subnets to your internal CIDR range (e.g. 10.10.0.0/24) and save.

— Named field and exact dashboard navigation with steps. changelog-20260908-46bc8c4b

05 **Cursor self-hosted machines integration** NEW 75

Adds Cursor self-hosted machines integration, running Cursor Cloud Agents on Cloudflare Containers with a Durable Object namespace, R2 bucket binding, and cron trigger via the open-source Cursor Cloudflare Workers template.

— Names components used but no setup steps given. changelog-20260908-46bc8c4b

06 **Bulk creation of Tunnel and Mesh routes** NEW 65

Supports bulk creation of Cloudflare Tunnel and Cloudflare Mesh routes in a single action using comma-separated CIDR ranges or hostnames, with Add another staging

and automatic removal of successfully created routes if the batch partially fails.

— Describes mechanism and failure handling but is UI-only. changelog-20260908-46bc8c4b

07 **Custom DHCP options in Appliance dashboard** NEW 60

Enables configuring custom DHCP options (PXE/iPXE boot, VoIP provisioning, vendor-specific, or arbitrary option code/type/value) from the Cloudflare One Appliance dashboard under LAN configuration > DHCP server options.

— Names option types and UI path but no walkthrough. changelog-20260908-46bc8c4b

08 **Log Explorer dataset configuration refresh** IMPROVED 60

Refreshes Log Explorer dataset configuration with grouped field selection (with per-category counts), field data-type details, bulk select/reset controls, and ingestion filters, accessible via Add dataset or Actions > Edit.

— Names UI entry points but lacks field-level specifics. changelog-20260908-46bc8c4b

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 8, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue