THE CYBER TOOLCHAIN

NO. 057

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 26, 2026 · EVERY DAY

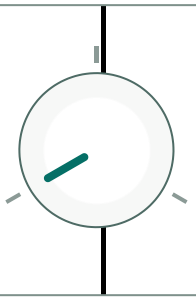
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 32 tools

OFFENSIVE SECURITY

vim ? help

◆ Password & Credential Cracking

brutespray

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

Fast, multi-protocol credential brute-forcer. Parses Nmap, Nessus, and Nexpose output to automatically test default and custom credentials across 30+ protocols.

> brutespray v2.7.2 turns the tool into an engagement-aware spraying platform, adding deterministic dry-run planning with acknowledgment gating, HMAC-chained audit verification, lockout-aware scheduling, and manifest-driven scope and evidence controls, alongside a new workspace store, HTTP auth templates, and a canonical service registry.

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

01 **Dry-run planning and acknowledged execution** NEW 92

`--dry-run` resolves targets and credentials into a deterministic JSON plan without any authentication attempts, `--plan-out` writes that plan to a file for review, approval, or CI evidence, and `--require-plan-ack` requires the exact reviewed plan hash to match before execution proceeds; the hash is bound to credential contents via keyed HMAC so same-sized credential-list substitutions invalidate the plan. Using `--require-plan-ack` requires an engagement manifest with `evidence.hmac_key` set, and runs without it fail closed.

Generate and save a reviewable execution plan from an Nmap scan before any credentials are attempted — useful for change-control or CI gating.

```
$ brutespray -f nmap.gnmap --dry-run --plan-out engagement-plan.json
```

Run the actual spray only after the plan hash has been reviewed and acknowledged, preventing credential substitution attacks between review and execution.

```
$ brutespray -f nmap.gnmap --require-plan-ack <plan-hash>
```

— Named flags with runnable before/after commands and HMAC mechanism. v2.7.2

02 **Authenticated audit chains and verification** NEW 78

Adds a `brutespray audit verify` subcommand that requires the `BRUTESPRAY_AUDIT_HMAC_KEY` environment variable, failing closed with a bad exit code if the key is missing or incorrect. Audit trails now use authenticated HMAC-SHA256 chaining instead of a recomputable plain hash, the manifest gains a required

`evidence.hmac_key` field for acknowledged execution and chain verification, and sensitive artifacts (checkpoints, session logs, stores, audit files) are written with owner-only permissions.

— Named subcommand and env var but no worked example. v2.7.2

03 **Engagement manifests with scope controls** NEW 76

New YAML engagement manifests carry engagement identity, authorization references, scope, lockout policy, and evidence defaults. Allow/deny scope controls add host and CIDR allowlists plus explicit deny rules applied before execution, with rejected targets appearing in dry-run output alongside stable reasons.

Define lockout-safe spraying policy and evidence mode in an engagement manifest so every run inherits the same budget and audit settings.

yaml

```yaml
lockout_threshold: 5
lockout_window: 30m
safe_margin: 1
jitter_percent: 10
evidence:
  mode: hash
  hmac_key: "<your-hmac-key>"
```

— Runnable manifest example shows the structure directly. v2.7.2

04 **Lockout-aware attempt budgets** IMPROVED 72

Engagement manifests can set `lockout_threshold`, `lockout_window`, `safe_margin`, and `jitter_percent` to control attempt budgets, with scheduling now shared across worker pools instead of being reset per goroutine. Account identities used for budgeting are normalized to `service|domain|user` across worker pools.

— Names all four policy keys and the before/after scheduling fix. v2.7.2

05 **Token-aware InfluxDB v2 execution** IMPROVED 70

Token-aware execution for InfluxDB v2 attempts each token once instead of multiplying it across default usernames; use `-m mode:v1` to preserve the previous username/password pair generation behavior.

— Names the exact flag and before/after attempt behavior. v2.7.2

THINNER COVERAGE BELOW

06 **Evidence modes for credential output** NEW 56

Adds evidence modes `full`, `redacted`, `hash`, and reserved `encrypted` for machine-readable credential output; `hash` mode emits HMAC-SHA256 correlation

values instead of printing the secret.

— Names all four modes but no dedicated example. v2.7.2

07 **Canonical service descriptor registry** NEW 46

Adds a canonical service descriptor registry covering default ports, aliases, credential modes, routing support, module parameters, stability, and wordlist aliases.

— Lists many covered attributes but no usage path shown. v2.7.2

08 **`-P` becomes strict print-and-exit** BREAKING 42

`-P` is now a strict print-and-exit operation in the shared `Execute` path; any setup that relied on `-P` continuing into credential attempts will no longer do so.

— Clear before/after behavior change but no migration example. v2.7.2

09 **JSONL workspace store** NEW 40

Adds a JSONL workspace store for durable engagement records that requires no CGO or external database, along with new workspace command paths for initialization and inspection.

— Describes the store but no exact command names given. v2.7.2

10 **Declarative HTTP auth templates** NEW 38

Adds declarative YAML templates describing HTTP/API authentication flows without enabling arbitrary shell execution; invalid or unsafe templates fail validation before any network execution.

— Explains safety behavior but no sample template or flag. v2.7.2

└──▷ **BREAKING ON UPGRADE**

! `--require-plan-ack` requires an engagement manifest with `evidence.hmac_key` set; executions using plan acknowledgment without this key will fail closed.

! `brutespray audit verify` requires the `BRUTESPRAY_AUDIT_HMAC_KEY` environment variable; missing or incorrect keys cause verification to fail.

! `-P` is now a strict print-and-exit operation in the shared `Execute` path; any setup that relied on `-P` continuing into credential attempts will no longer do so.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Snyk

Feedback ⓘ 1 RELEASE · 2026-08-26 DOCS ↗ BLOG ↗

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

> Snyk Evo now scopes repository, asset, and issue visibility to a user's Snyk Organizations via the existing role hierarchy, changing what the Tenant Viewer and Tenant Member roles grant by default.

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

01 **Organization-scoped access in Evo** BREAKING 75

The Tenant Viewer role now grants scoped read-only access to assets and issues limited to the Organizations the user belongs to, replacing the previous full-Tenant visibility. The Tenant Member role now carries no Evo access at all, letting administrators withhold Evo without removing a user's broader Snyk platform access, and new users default to Tenant Member on join, requiring an explicit role change to gain any Evo access. Organization names now appear alongside repositories in Evo to disambiguate the same repository imported into multiple Organizations, and assets without an assigned Organization — currently all Agent Supply Chain Security and Continuous Offensive Security data — are hidden from scoped users.

— Names exact roles and behavior changes but no config key or command given.

2026-08-26-c1039ba1

WAS THIS USEFUL? 👍 👎

Aikido Security

Feedback ⓘ 4 RELEASES · SEEN 2026-08-26

NOTES ↗

Aikido Security is a developer-focused security platform that scans code, infrastructure, and dependencies for vulnerabilities and misconfigurations.

> Aikido added pentest-grade LLM-driven code auditing with one-click AutoFix, a Data Exposure Audit that traces sensitive data flows without connecting data stores, LLM agents that auto-triage secrets and assess CVE exploitability to cut noise, and a blocked-install approval workflow to guard company devices against supply chain attacks.

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

01 **Data Exposure Audit for sensitive data flows** NEW 65

A new Data Exposure Audit uses LLM agents to trace sensitive data from sources (request bodies, uploads, DB reads, config, app logic) into sinks (logs, third-party APIs, LLM prompts, caches, queues, storage), working entirely from already-connected repos without needing to connect data stores. Each audit consumes credits from the Aikido Wallet.

– Detailed mechanism and named sources/sinks but no usage steps given

🚀DSPM: Find data exposure risks withou…

02 **Blocked-install approval workflow and team policies** NEW 65

A request-and-approval workflow pauses blocked package installs and routes a one-click exception request to the admin Inbox inside Aikido, where admins can review and approve or deny it. Group-based policies let admins set different install rules per team.

Review and act on a developer's blocked-install exception request from the admin Inbox.

> 📍 In the Aikido console, go to Inbox, find the pending install request, review the flagged package name and policy reason, then approve or deny the exception.

– Exact admin UI flow given but policy config options not detailed

🚀Protect Company Devices From Supply C…

03 **Pentest-grade Code Audit without live environment** NEW

60

Code Audit surfaces code-based evidence, root cause analysis, and remediation guidance for a codebase with no environment required, plus AI-powered AutoFix that creates remediation PRs in a few clicks. Pricing is shown upfront in credits, calculated from the size and complexity of the selected codebases, before the audit starts.

Start a pentest-grade code audit on a repo with no live environment — useful for auditing a smart contract or mobile app codebase.

> 📍 In the Aikido dashboard, go to Code Audit, connect the target repository, review the upfront credit cost, then click Start Audit.

– Explains flow and pricing but not audit mechanics or accuracy

🚀Code Audit: Pentest-grade reasoning o…

---

04 **Secrets AutoTriage for automated secret validation** NEW

60

Secrets AutoTriage runs an LLM agent after entropy detection to assess whether a discovered secret is a real, exposed credential, determine its liveness and scope, and automatically ignore or escalate findings based on user-defined rules. Each run is logged with a short rationale so teams can review what was analyzed and what action was taken.

– Clear mechanism described but no rule-config surface or steps shown

🚀Extreme Noise Reduction for Secrets a…

---

05 **CVE Exploitability Analysis with automated remediation** NEW

60

An LLM agent evaluates whether a CVE is practically exploitable given how the codebase actually uses the affected package, then can automatically upgrade, downgrade, snooze, or ignore the finding. Runs are logged with a short rationale, and remediation actions are driven by user-controlled rules combined with codebase-aware reasoning.

– Mechanism and possible actions are named but no rule-setup details

🚀Extreme Noise Reduction for Secrets a…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ 1 RELEASE · SEEN 2026-08-26 NOTES ↗

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

> Semgrep's latest snapshot adds org-wide nosemgrep controls in public beta, a Workflow issues column on the Projects page, and tightens Supply Chain diff-scan license checks, alongside two breaking changes tied to Unified Policies migration and Multimodal access.

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

01 **Policies V1 and Ruleboard endpoints return 410 Gone** BREAKING

70

For organizations migrated to Unified Policies, the Policies V1 and Ruleboard endpoints now return `410 Gone` with a `DEPLOYMENT_MIGRATED` code instead of a normal response.

— Exact status code and error code named, directly actionable snapshot-20260826

02 **Supply Chain diff-scan license checks without dependency changes**

IMPROVED

68

Supply Chain now runs license checks during diff scans even when no dependency source file (lockfile or manifest) changed in the merge request; this requires Semgrep CLI 1.174.0 or later.

— Names exact CLI version requirement, a clear upgrade action snapshot-20260826

03 **Org-wide nosemgrep setting in public beta** NEW

60

Introduces an organization-wide `nosemgrep` setting (public beta) that lets admins enable or disable `nosemgrep` comments org-wide, and lets users filter findings by whether they were ignored via a `nosemgrep` comment.

— Named setting and behaviour, but no exact config path given snapshot-20260826

THINNER COVERAGE BELOW

04 **403 responses when Multimodal is disabled** BREAKING

53

Organizations without Semgrep Multimodal enabled now receive a `403` response from every trigger path, including scheduled and API-initiated runs.

— Names exact response code and affected trigger paths snapshot-20260826

05 **Workflow issues column on Projects page** NEW 50

The Projects page gains a Workflow issues column; selecting a count opens the Agentic Workflows issues for that project.

— Names the UI surface but no further mechanism given snapshot-20260826

06 **Multimodal defaults to GPT-5.6 Luna** IMPROVED 48

Semgrep Multimodal now defaults to GPT-5.6 Luna in more places, with clearer validation when adding OpenAI credentials via bring-your-own-key; alternative models remain available for bring-your-own-key and fallback.

— Names model and BYOK flow but no config steps snapshot-20260826

07 **Faster target selection for Supply Chain incident scans**

IMPROVED 23

Improves performance of target selection for Supply Chain incident scans on organizations with large numbers of projects.

— Vague performance claim with no numbers or mechanism snapshot-20260826

└──▷ **BREAKING ON UPGRADE**

! For organizations migrated to Unified Policies, the Policies V1 and Ruleboard endpoints now return `410 Gone` with a `DEPLOYMENT_MIGRATED` code instead of a normal response.

! Organizations without Semgrep Multimodal enabled now receive a `403` response from every trigger path, including scheduled and API-initiated runs.

WAS THIS USEFUL?  

Socket

Feedback

3 RELEASES · 2026-08-20 → 2026-08-25

BLOG ↗

Socket detects supply chain attacks by analyzing open source package dependencies for malicious code, vulnerabilities, and suspicious behaviors.

> Socket added Asana integration for alert ticketing, moved PHP and Composer support to Beta with reachability analysis now generally available, and introduced experimental Firefox extension scanning for Enterprise customers.

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

01 **Asana integration for alert ticketing** NEW 65

A new beta Asana integration (Business and Enterprise plans) lets teams create and manage Asana tasks directly from Socket alerts, either manually or via ticketing rules. Ticketing rules can filter alerts by event, type, priority, repository and other attributes, with per-rule configuration of Asana workspace, project, tags and assignee, plus bi-directional status sync between linked tasks and alerts.

— Names plan tiers and config fields but no API/CLI surface given.

New: Asana integration for alert ticket…

THINNER COVERAGE BELOW

02 **Firefox extension protection** NEW 56

Socket now scans every extension listed in Mozilla's official add-ons directory, monitoring new releases for changes in permissions, code, network activity, and behavior. It detects malware, credential and clipboard theft, data exfiltration, remote loading, obfuscation, impersonation, and related campaign activity. This coverage is available in Experimental for Enterprise customers, who must contact their Socket account team to enable it.

— Requires contacting account team; no self-serve config or endpoint named.

New: Firefox extension protection (2026…

03 **PHP and Composer support graduates to Beta with GA reachability**

IMPROVED

52

PHP and Composer support moved from Experimental to Beta and is now available to all Socket users. PHP reachability analysis is now generally available for Team plan customers and above with no Experimental access required, including precomputed reachability and full application reachability.

— States plan and feature scope but no config or command to invoke.

PHP and Composer support is now in beta…

WAS THIS USEFUL? 👍 👎

GitHub

Feedback

4 RELEASES · 2026-08-25

BLOG ↗

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

> GitHub shipped GA rule insights dashboards with CSV export, path-level exceptions for push rules, a consolidated Customize tab in the Copilot app, and a way to block abusive users straight from security advisories.

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

01 Path exceptions for push rules in rulesets IMPROVED

90

Adds path exception patterns to the Restrict file paths push rule, letting you block a file type repository-wide while exempting specific glob patterns like `**/gradle/wrapper/*.jar`, and to the Restrict file size push rule, letting you enforce a size limit without requiring immediate changes to pre-existing oversized files at exempted paths. GitHub validates exception path patterns at ruleset save time, surfacing pattern errors immediately.

Allow JAR files only inside the Gradle wrapper directory while blocking them everywhere else in the repository.

> 📍 In your repository or organization, go to Settings › Rules › Rulesets, open or create a ruleset, add the Restrict file paths push rule, expand Allowed exceptions, and add the pattern `**/gradle/wrapper/*.jar`.

Adopt a repository-wide file size limit without forcing an immediate cleanup of existing large files at known legacy paths.

> 📍 In your repository or organization, go to Settings › Rules › Rulesets, open or create a ruleset, add the Restrict file size push rule, expand Allowed exceptions, and add the path patterns for directories that already contain files exceeding the limit.

— Names exact rule types, glob pattern, and full UI navigation to configure.

Push rules in rulesets now support path…

02 Customize tab in GitHub Copilot app NEW

75

New Customize tab in the GitHub Copilot app consolidates MCP servers, plugins, skills, and canvases into a single discovery surface with featured picks and category browsing. Includes trending MCP server discovery by category and featured canvases that bridge understanding and action, such as Azure DevOps backlog triage and prioritization delegated to Copilot.

Reach the new Customize tab to discover and install MCP servers, plugins, skills, and canvases for your team.

> 📍 1. Open the GitHub Copilot app. 2. Select `Customize` in the top navigation. 3. Browse `Featured` for curated picks, or choose a dedicated section (MCP servers, plugins, skills, canvases) to filter by type. 4. Select a trending MCP server or featured canvas and follow the install prompt.

— Includes exact navigation steps to reach and use the feature.

GitHub Copilot app Customize tab is gen…

---

03 **Rule insights dashboard GA at repo and org levels** NEW

70

Adds an organization-level rule insights dashboard at Settings > Repository tab, showing aggregated ruleset evaluation metrics, top repositories by bypass count, and filters by evaluation status, branch, ruleset, and date range across all repositories. Adds a repository-level dashboard at Settings > Rules tab showing successes, failures, and bypasses over time plus most active bypassers, with charts linking to the insights page with filters prefilled. Both levels support CSV export of rule insights data for offline reporting and record keeping.

— Names exact settings tabs and metrics but no CLI/API path.

Rule insights dashboard generally avail…

---

04 **Block users from security advisories** NEW

70

Adds a block-user action to the three-dot menu on security advisory descriptions and comments in public repositories, eliminating the need to navigate to organization or profile settings. Shows a confirmation dialog explaining the consequences before the action is taken, and restricts the in-advisory block action to organization moderators or administrators for org-owned repos, or the account owner for personal repos.

Block a spammer directly from a security advisory comment without leaving the page.

> 📍 In the security advisory, open the three-dot menu on the offending description or comment, then select `Block user` and confirm the action in the dialog that appears.

— Clear steps and permission scope given, no API surface.

Block users directly from security advi…

WAS THIS USEFUL? 👍 👎

sbomify

Feedback

1 RELEASE · 2026-08-25

NOTES ↗

sbomify is a product security artifact hub and a trust center.

> sbomify v26.8.0 enriches vulnerability findings with malicious-package detection, EPSS/CVSS 4.0 scoring and ENISA EUVD flagging, adds formal risk acceptance and patch SLA/CRA KPI tracking, and redesigns the Trust Center with per-product advisory gating; smaller updates touch access tokens, CI/CD integration, Settings and compliance exports.

└──▷ **WHAT SHIPPED · 21 FEATURES** most completely described first ? what's the number?

01 **Trust Center redesign with per-product advisory gating**

IMPROVED 85

Redesigns the Trust Center with per-product advisory gating for published advisories, adds a searchable products table, and cleans up the products drill-down. To publish: open Advisories > New Advisory, select affected products and version range in the modal, then toggle `Publish to Trust Center` per product before saving.

Publish a security advisory to the Trust Center with per-product gating, so only relevant products expose the advisory to external consumers.

> 📍 In the sbomify UI, open Advisories › New Advisory, select the affected products and version range in the modal, then toggle `Publish to Trust Center` per product before saving.

— Named gating mechanism plus exact publish steps v26.8.0

02 **Formal risk acceptance for vulnerabilities** NEW 80

Adds formal risk acceptance for vulnerabilities, recording an owner and an expiry date. On a vulnerability's detail page, open the `Risk Acceptance` panel, enter the owner name and expiry date, then click `Accept` .

Record formal risk acceptance for a vulnerability finding — capturing an owner and expiry — so the risk appears tracked rather than ignored in compliance reports.

> 📍 On a vulnerability's detail page, open the `Risk Acceptance` panel, enter the owner name and expiry date, then click `Accept` .

— Exact UI steps and captured fields given v26.8.0

03 **On-demand plugin assessment re-runs** NEW 70

Enables on-demand assessment re-runs for plugins without waiting for a scheduled scan. Open the artifact's Plugin Assessments panel and click `Re-run` next to the assessment you want to refresh.

Trigger an on-demand plugin assessment re-run after uploading a new SBOM, without waiting for the next scheduled scan.

> 📍 In the sbomify UI, open the artifact's Plugin Assessments panel and click `Re-run` next to the assessment you want to refresh.

— Exact UI steps to trigger the re-run `v26.8.0`

---

04 **Vulnerability finding enrichment: EPSS, CVSS 4.0, malicious packages, ENISA EUVD**

NEW 60

Captures EPSS scores and CVSS 4.0 from both scanner paths, surfaces malicious packages as their own distinct vulnerability class separate from CVE-based findings, and flags ENISA EUVD exploited findings alongside existing CISA KEV badges on vulnerability views.

— Names specific scoring feeds and flags but no navigation given `v26.8.0`

THINNER COVERAGE BELOW

05 **Vulnerability posture metrics: patch SLA, CRA KPIs, lifecycle age/MTTR**

NEW 55

Adds a patch SLA matrix and five CRA KPIs to the vulnerability posture view, plus vulnerability lifecycle tracking for age and Mean Time To Remediate (MTTR).

— Named metrics but no view path `v26.8.0`

06 **New Advisory modal mapped to products and releases** NEW

55

Maps the New Advisory modal onto products and releases, with version selection and range-picking. In Advisories > New Advisory, select the affected products and version range in the modal before saving.

Publish a security advisory to the Trust Center with per-product gating, so only relevant products expose the advisory to external consumers.

> 📍 In the sbomify UI, open Advisories › New Advisory, select the affected products and version range in the modal, then toggle `Publish to Trust Center` per product before saving.

— Exact navigation given, mechanism partly described `v26.8.0`

07 **CI/CD integration dialog with embedded token** IMPROVED

55

Consolidates the CI/CD integration dialog into a single command with a 7-day token embedded.

— Names the token lifetime but not the command itself v26.8.0

08 **Settings reorganization into tabbed section pages** IMPROVED

50

Reorganizes Settings into one page per section with tabs across the top, and consolidates plugin access to a single entry point in Settings.

— Names the Settings area as the starting point v26.8.0

09 **Vulnerability drill-down filtering and result states** IMPROVED

45

Adds a suppressed-findings toggle plus state and KEV filters on the internal vulnerability drill-down, and distinguishes nothing scanned from nothing found in vulnerability results.

— Names the filters but not their exact location v26.8.0

10 **Access token expiry controls** IMPROVED

45

Adds hour-scale expiry presets and an arbitrary custom TTL for access tokens, and warns users before access tokens expire.

— Names the preset/TTL mechanism but no exact UI path v26.8.0

11 **EU Authorized Representative field on Declaration of Conformity**

NEW

45

Captures the EU Authorized Representative field and includes it on the Declaration of Conformity.

— Names the field and document it appears on v26.8.0

12 **End-of-life transition readiness tracking** NEW

45

Adds end-of-life transition readiness tracking, announcement, and final artifact handling.

— Describes stages but no concrete UI or workflow v26.8.0

13 **Risk categories for SPDX license identifiers** NEW

40

Adds risk categories for SPDX license identifiers in the licensing view.

— Names the licensing view but not the categories v26.8.0

14 **Supplier list entity for teams** NEW 40

Adds a supplier list entity to teams for collecting artifacts from external suppliers.

— Names the entity but no schema or workflow detail v26.8.0

15 **CPE identifier validation and copyable values** IMPROVED 40

Validates CPE identifiers and makes identifier values copyable in the UI.

— Names the change but not validation rules v26.8.0

16 **Art. 13(15) retention floor on compliance exports** NEW 40

Pins the Art. 13(15) retention floor on compliance export packages.

— Names the regulation clause, no export path given v26.8.0

17 **OpenChain Telco SBOM Guide conformance plugin** NEW 35

Adds an OpenChain Telco SBOM Guide v1.1 conformance plugin.

— Names the standard and version, nothing more v26.8.0

18 **SBOM freshness window and stale flag** NEW 35

Introduces an SBOM freshness window and stale flag to signal outdated SBOMs.

— Names the mechanism but no thresholds or config v26.8.0

19 **Per-IP rate limiting for anonymous and public endpoints** NEW 35

Adds per-IP rate limiting for anonymous and public endpoints.

— No rate limits or endpoint names given v26.8.0

20 **Navbar search as navigation spotlight** IMPROVED 30

Turns the navbar search into a navigation spotlight.

— Bare description of a UI change v26.8.0

21 **Advisory IDs linked to upstream database** IMPROVED 25

Links advisory IDs directly to the upstream database that holds them.

— Thin one-line mention, no source database named v26.8.0

WAS THIS USEFUL? 👍 👎

npq

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗ CODE ↗

safely install npm packages by auditing them pre-install stage

> npq v3.27.0 adds explicit opt-in support for running its pre-install audit safely inside CI and other non-interactive environments.

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

01 **Non-interactive install authorization for CI** NEW 80

Adds the `--allow-non-interactive-install` CLI flag and the `NPQ_ALLOW_NON_INTERACTIVE_INSTALL=true` environment variable, either of which explicitly authorizes npq to proceed with installs in non-interactive (CI/automated) environments instead of blocking them.

Authorize a package install in a headless CI pipeline where no TTY is present, so npq proceeds rather than blocking.

```
$ NPQ_ALLOW_NON_INTERACTIVE_INSTALL=true npq install express
```

Use the CLI flag directly in a CI script step to allow non-interactive routing without setting environment variables.

```
$ npq install express --allow-non-interactive-install
```

— Names exact flag and env var with runnable commands. v3.27.0

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

sstimap

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

Automatic SSTI detection tool with interactive interface

> SSTImap v1.4 adds session persistence for detected vulnerabilities, three new template engine detection plugins, and authenticated crawling that respects user-supplied headers and cookies.

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

01 **New template engine detection plugins** NEW 55

New plugins add detection and exploitation for the Mustache template engine's GHSA-5hrj-7fc3-cmx9 vulnerability, Pybars3 and Pybars4 template engines, and a fuzzer for uncommon and legacy evaluating template injection patterns.

— Names specific engines and CVE but no usage steps v1.4

02 **Save and reload detected vulnerabilities** NEW 35

SSTImap can now save detected vulnerabilities to a file and reload them in a later session, avoiding re-scanning.

— Describes capability but no file format or flag named v1.4

03 **Header and cookie-aware crawling** IMPROVED 35

Crawling and form detection now respect user-supplied headers and cookies, enabling authenticated crawl sessions.

— States behaviour change but no flag or config named v1.4

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

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 System Recovery 8.38 bakes BitLocker TPM exploit provisioning directly into the boot media installer, adds checksum verification to disk imaging, and introduces browser artefact extraction.

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

01 **Built-in BitLocker TPM exploit provisioning** NEW 85

During Elcomsoft System Recovery installation, when prompted, you can insert a second USB flash drive and select `YellowKey` or `GreatXML` to provision a BitLocker TPM exploit drive alongside the main boot drive, letting examiners unlock TPM-protected encrypted drives on-site without the recovery key.

During ESR installation, provision a second USB drive loaded with the YellowKey or GreatXML BitLocker TPM exploit so you arrive on-site ready to unlock encrypted drives without the recovery key.

> 📍 During Elcomsoft System Recovery installation, when prompted, insert a second USB flash drive and select YellowKey or GreatXML to provision the BitLocker exploit drive alongside your main boot drive.

– Names exact installer options and exploit names with steps

Elcomsoft System Recovery 8.38: Built-I…

THINNER COVERAGE BELOW

02 **Checksum logging for disk imaging** NEW 25

Adds checksum logging for disk images during the imaging workflow, alongside updates to the imaging library.

– No mechanism or checksum algorithm specified Elcomsoft System Recovery 8.38: Built-I…

03 **Browser artefact extraction** NEW 10

Elcomsoft System Recovery 8.38 adds browser artefact extraction capability, per the release headline.

– Mentioned only in headline, no detail given Elcomsoft System Recovery 8.38: Built-I…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

Fastly is a content delivery network and edge computing platform that accelerates website and application performance globally.

> Fastly's docs and Rust SDK update this window centered on new edge security/observability products—API Enforcement and Log Explorer & Insights—plus a batch of Rust SDK 0.13.1 additions covering caching, bot detection, and backend healthchecks.

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

01 **Rust SDK 0.13.1 additions** NEW 75

The Rust SDK adds `Request::as_handles`, `from_handle` constructors for `ConfigStore`, `SecretStore`, and `Log`, and a `ToBackend` implementation for `Arc<Backend>`. It also adds cache lookup timeout support respected for HTTP guest caching, `BotCategory::Headless` to the list of known bot categories, and experimental support for dynamic backend healthchecks.

– Names every new SDK method and constant but no usage example

Rust SDK 0.13.1 (2026-08-24)

THINNER COVERAGE BELOW

02 **API Enforcement for schema-based request validation** NEW 45

A new API Enforcement product validates incoming API requests against user-defined schemas and logs or blocks non-conforming requests at the Fastly edge.

– Names the product and behavior but no config or endpoint detail product docs

03 **Log Explorer & Insights for on-platform log monitoring** NEW 30

A new Log Explorer & Insights feature enables proactive storage, inspection, and monitoring of log data directly on the Fastly platform.

– High-level description only, no mechanism or UI path given product docs

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Chainguard

Feedback ⓘ 1 RELEASE · 2026-08-26 CHANGELOG ↗

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

> Chainguard's biggest change this window is the general availability of API v2 at `/v2/` (replacing `v2beta1`), alongside early-access SCIM user provisioning, self-service Helm chart provisioning in the Console, and 27 new container images including several FIPS variants.

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

01 **SCIM user provisioning for Okta and Entra ID** NEW 68

Supports SCIM user provisioning (Early Access) via the `/Users` resource for Okta and Microsoft Entra ID; assigning a user creates or updates their Chainguard provisioning record, and deactivating or unassigning them blocks the next login or token refresh within approximately one hour.

— Names endpoint and identity providers but early access with no full setup steps

changelog-20260826-614a0a18

THINNER COVERAGE BELOW

02 **27 new container images including FIPS variants** NEW 48

Adds 27 new container images this week, including FIPS variants for Airbyte components (`airbyte-bootloader`, `airbyte-cron`, `airbyte-worker`, `airbyte-workload-api-server`, `airbyte-workload-launcher`), `percona-server-mongodb-fips`, `commercial-nginx-plus-fips`, `apache-activemq-artemis-fips`, and `kubernetes-csi-external-health-monitor-fips`, among others.

— Lists concrete image names but no usage or config detail changelog-20260826-614a0a18

└──▷ **BREAKING ON UPGRADE**

! Existing `v2beta1` integrations must update their path prefix from `/v2beta1/` to `/v2/`; request and response shapes are preserved.

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Alibaba Cloud OpenSandbox

Feedback ⓘ

1 RELEASE · 2026-08-26 NOTES ↗

Secure, Fast, and Extensible Sandbox runtime for AI agents.

> OpenSandbox v0.1.16 adds resilient retry handling to its clients, lifecycle hooks and detached background runs for sandboxes, richer Pool capacity visibility and back-pressure signaling, and a major speedup in idle sandbox cleanup.

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

01 Retry transport with configurable RetryPolicy NEW

80

Async and sync clients gain configurable retry transport with backoff, jitter, Retry-After handling, per-attempt timeouts, and an overall deadline. Retries are on by default for idempotent requests on HTTP 429, 502, and 503; streaming SSE requests bypass replay; call RetryPolicy.disabled() to restore the previous fast-fail behavior.

— Names exact status codes and opt-out method, no runnable snippet python/sandbox/v0.1.16

02 Sandbox lifecycle hooks NEW

76

Adds `preStart` and periodic lifecycle hook parameters to `Sandbox.create` and `SandboxSync.create`, with server-authoritative timeout validation.

Register a preStart hook when creating a sandbox to run initialisation logic before the sandbox becomes ready.

python

```python
from opensandbox import Sandbox

sandbox = Sandbox.create(
    preStart={"command": ["./init.sh"], "timeout": 30}
)
```

— Named API methods plus a runnable code example python/sandbox/v0.1.16

03 Pool capacity visibility and back-pressure signaling NEW

74

Sandbox responses now expose confirmed Pool allocation summaries and add execd init mode and hardening-layer state fields to isolated capabilities, absent when the server cannot confirm the state. Exhausted Kubernetes Pool capacity now surfaces as HTTP 429 with error code `KUBERNETES::POOL_CAPACITY_EXHAUSTED` and a Retry-After

header instead of a generic readiness timeout; sandbox creation is not automatically replayed on this status.

— Names exact error code and header, no direct code sample `python/sandbox/v0.1.16`

THINNER COVERAGE BELOW

04 **Faster idle sandbox pool cleanup** IMPROVED 48

Bounded concurrency for idle sandbox pool cleanup cuts the release time for 100 sandboxes from approximately 8.3 seconds to 0.8 seconds.

— Concrete before/after numbers but no user-facing action `python/sandbox/v0.1.16`

WAS THIS USEFUL?  

DETECTION & RESPONSE

◆ SOAR

Tines

Feedback ⓘ 1 RELEASE · 2026-08-26 CHANGELOG ↗

Tines lets every team build. IT and Security maintain complete visibility, governance, and control.

> Tines 43.0.0 brings native Tines primitives into Apps and introduces Workbench for Cases, a dedicated environment for working with case data.

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

01 **Tines primitives support in Apps** NEW 35

Apps can now use native Tines building blocks (primitives) within app configurations, enabling use of these blocks directly inside apps.

— Names the capability but not which primitives or mechanism. changelog-20260826-f0faad44

02 **Workbench for Cases** NEW 35

A new environment called Workbench for Cases lets teams work with case data.

— Describes the new environment but not its features or workflow.

changelog-20260826-f0faad44

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Feedly Threat Intelligence

Feedback ⓘ

1 RELEASE · 2026-08-25 BLOG ↗

Collect, analyze, share actionable intel faster - Feedly TI

> Feedly Threat Intelligence shipped a new Splunk Hunting Queries AI Model that auto-collects detection rules across six formats from 5,000 vetted sources, plus surfaces those queries directly on CVE Insights Cards.

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

01 Splunk Hunting Queries AI Model for detection rules NEW

78

A new AI Model continuously collects and extracts hunting queries tailored to configured threat intelligence requirements from vendor blogs, threat reports, GitHub repos, and researcher write-ups across 5,000 vetted CTI sources. It recognizes six detection rule formats — Splunk SPL, Sigma, YARA, Suricata, Snort, and KQL (Microsoft Defender and Sentinel) — covering network traffic inspection, file and malware matching, and log-based detection. Multiple AI Models can be paired in an AI Feed to narrow collected queries to specific topics, threats, or malware families, and results can be retrieved programmatically via the API for direct integration into a detection pipeline.

· · ·

— Names sources, six formats and API, but no endpoint or flag.

Automatically collect Splunk Hunting Qu…

THINNER COVERAGE BELOW

02 CVE Insights Card surfaces Splunk Hunting Queries NEW

55

The CVE Insights Card now surfaces any Splunk Hunting Query associated with a CVE, letting users download the query and jump to the source article for surrounding context directly from the card.

— Describes UI location and action but no config or API detail.

Automatically collect Splunk Hunting Qu…

└──▷ **ALSO FROM THESE RELEASES**

· · ·

WAS THIS USEFUL?  

MISP

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

MISP (core software) - Open Source Threat Intelligence and Sharing Platform

> MISP v2.5.45 overhauls LDAP authentication with nested group support and proxy-header login, adds a public `/gpg.asc` key endpoint and CLI worker management, and completes the Bootstrap 5 Overmind UI migration across dozens of admin views alongside several ACL enforcement fixes.

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

01 LDAP nested groups, role/org mapping, and proxy auth NEW

85

Adds `ldapNestedGroups` to resolve nested LDAP group membership for role and org mapping, plus `ldapOrgGroupMapping` and `ldapRoleGroupMapping` settings (precedence follows settings order) supporting mapping via directory field or group membership. Adds LDAP sync of MISP accounts via the `User check_validity` admin task, authentication against a header asserted by a trusted proxy, and a setting to disable MISP accounts flagged `ACCOUNTDISABLE` in `userAccountControl`. Adds a pytest suite under `tests/ldapauth/` running against an OpenLDAP service in CI for LDAP authentication coverage.

— Names every config key and task but no full walkthrough v2.5.45

02 Instance GPG key endpoint NEW

75

Serves the instance GPG key at `/gpg.asc` with caching, fetchable via `curl https://<misp-instance>/gpg.asc` — useful for verifying signed exports or configuring encrypted sharing.

Fetch the instance public GPG key from a MISP server — useful for verifying signed exports or configuring encrypted sharing.

```
$ curl https://<misp-instance>/gpg.asc
```

— Exact endpoint plus a runnable curl command v2.5.45

03 Access control enforcement fixes IMPROVED

75

Enforces per-org module restrictions on all getEnabledModule() paths, enforces the Event modify ACL on attribute enrichment requests, and enforces a `distribution=0` filter in REST search `fetchEvent`.

— Names the exact enforcement points and method v2.5.45

04 **Bootstrap 5 Overmind UI migration** IMPROVED 70

Migrates Server settings, Workflow indexes (ad-hoc, triggers, modules), the workflow editor, Decaying Model index, misp-module enrichment, and registration/user settings/password flows to the Bootstrap 5 Overmind UI. Adds new Bootstrap 5 views for Jobs, Tasks, AnalystData, Feeds, TaxiiServers, CryptographicKeys, and GalaxyClusterRelations.

— Lists every migrated area but gives no behavioural change detail v2.5.45

THINNER COVERAGE BELOW

05 **Import UX additions across event, server and feed views** NEW 55

Adds an import action on the event index and server preview views, plus a freetext feed preview index and import modals; feed migration to the Bootstrap 5 (Overmind) UI is now complete.

— Names the views affected but not the exact UI flow v2.5.45

06 **CLI management of Supervisor-backed workers** NEW 45

Enables managing Supervisor-backed workers directly from the CLI.

— States the capability but no command syntax v2.5.45

07 **Organisation merge feature** NEW 30

Adds an organisation merge feature in the UI.

— Bare description, no mechanism or location given v2.5.45

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

OSSEC HIDS

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

OSSEC is an Open Source Host-based Intrusion Detection System that performs log analysis, file integrity checking, policy monitoring, rootkit detection, real-time alerting and active response.

> OSSEC 4.3.0 adds a maintenance mode to pause FIM alerting during planned changes, brings TCP/TLS transport and a larger payload limit to syslog forwarding, and adds ModSecurity audit log collection alongside a batch of new decoders, rules, and agent-management hardening.

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

01 **FIM maintenance mode via agent_control -M** NEW 91

`agent_control -M` gains subcommands `enable`, `disable`, `status`, and `end` (plus `-u <id>` to target a specific agent), letting syscheck absorb file changes into the baseline without firing integrity alerts; `end` restarts syscheck/rootcheck and clears maintenance after the scan completes.

Pause FIM alerting on a specific agent during a planned change window so file modifications are absorbed into the baseline without generating integrity alerts.

```
$ agent_control -M enable -u 042
# ... perform maintenance ...
agent_control -M end -u 042
```

— Full subcommand set, flag, and behaviour with runnable example. 4.3.0

02 **TCP/TLS syslog output with larger payload limit** NEW 90

`ossec-csyslogd` `syslog_output` gains TCP and optional TLS support via new config keys `protocol`, `tls`, `tls_verify`, and `tls_ca`, alongside existing UDP; alert payloads in default, CEF, JSON, and Splunk formats now support up to `OS_MAXSTR` instead of the former 2048-byte limit.

Forward OSSEC alerts to a remote syslog collector over TLS-secured TCP instead of UDP, with server certificate verification.

```xml
<syslog_output>
  <server>siem.example.com</server>
  <port>6514</port>
  <protocol>tcp</protocol>
  <tls>yes</tls>
  <tls_verify>yes</tls_verify>
  <tls_ca>/var/ossec/etc/ca.pem</tls_ca>
  <format>json</format>
</syslog_output>
```

— Named config keys, formats, and a before/after byte limit. 4.3.0

03 **ModSecurity audit log collection** NEW 72

Adds a `modsec-audit` `localfile` format to collect libmodsecurity serial audit logs, plus new decoders and rules for nginx ModSecurity error-log events.

Collect libmodsecurity serial audit logs from a web application firewall for decoding and alerting by OSSEC.

```xml
<localfile>
  <log_format>modsec-audit</log_format>
  <location>/var/log/modsec_audit.log</location>
</localfile>
```

— Named format and runnable config, thin on decoder detail. 4.3.0

04 **New decoders and rules for Postfix, Snort, BIND, cPanel, Windows and Asterisk**

NEW 60

Adds Postfix SMTP connect-abort detection with escalation by source IP; matches Snort `[Drop]` fast alerts in the Snort decoder; decodes named client IPs with an optional BIND `@cookie` prefix; adds cPanel/cpsrvd decoders and rules; adds Windows interactive logon (logon type 2) success detail rule `18262`; and adds an Asterisk IPv6 denied decoder so `srcip` omits square brackets.

— Each rule/decoder named but no usage path given. 4.3.0

05 **Agent management safeguards in manage_agents** IMPROVED 60

`manage_agents -f -` allows bulk-loading agents from stdin, agent state is backed up before remove/force-delete, and OSSEC now alerts on duplicated IP.

— Named flag and behaviour changes, no full example given. 4.3.0

THINNER COVERAGE BELOW

06 **Default Linux compile/link hardening** IMPROVED 40

Enables Linux compile/link hardening by default, including PIE, FORTIFY, stack protector, and full RELRO.

— Names hardening techniques but nothing for a reader to act on. 4.3.0

07 **Dedicated iptables chain for firewall-drop response** IMPROVED 35

Firewall-drop active response now uses a dedicated OSSEC iptables chain instead of the default chain.

— Names the mechanism change but no configuration detail. 4.3.0

08 **syscheck_control flag validation** IMPROVED 35

Rejects dangerous or ambiguous `syscheck_control` flag combinations such as `-r -u`.

— Names the flags but not the full set of rejected combinations. 4.3.0

09 **CDB list trailing comment support** NEW 28

Adds support for `###` trailing comments in CDB list text files.

— Bare, one-line description of a minor syntax addition. 4.3.0

WAS THIS USEFUL?  

Trend Vision One

Feedback ⓘ

2 RELEASES · 2026-08-24 → 2026-08-25 BLOG ↗

Trend Micro's unified security platform providing threat intelligence, endpoint protection, and attack surface visibility across cloud and on-premises environments.

> Trend Vision One rolled out generally available zero-day risk events sourced from Trend Micro's Zero Day Initiative and added click-through navigation from Cyber Risk Subindex asset totals into filtered asset views.

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

01 Clickable asset totals in Cyber Risk Overview IMPROVED

65

Clicking an asset type total in Cyber Risk Overview now opens Attack Surface Discovery filtered to that asset type and group, letting you drill into the specific assets contributing to a Cyber Risk Subindex score.

When you want to investigate which specific assets are driving a Cyber Risk Subindex score for a defined asset group.

> 📍 In Trend Vision One, go to Cyber Risk Exposure Management › Cyber Risk Overview, locate your asset group, and click an asset type total under the relevant Cyber Risk Subindex to open Attack Surface Discovery filtered to that asset type and group.

— Includes exact click-through navigation steps to reproduce.

Clickable asset totals allow you to vie…

02 ZDI-powered zero-day risk events (GA) NEW

60

TrendAI™ Zero Day Initiative (ZDI) `-powered` risk events are now generally available, surfacing zero-day vulnerabilities affecting software in your environment. They appear under Cyber Risk Exposure Management → Continuous Risk Management → Vulnerability Management and Cyber Risk Exposure Management → Continuous Risk Management → Threat and Exposure Management.

— Names exact navigation paths but no mechanism detail.

TrendAI™ Zero Day Initiative risk event…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Intruder

Feedback ⓘ changelog · Aug 26, 2026 CHANGELOG ↗

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

Adds Bitbucket Cloud as a supported source control integration for AI Pentesting, enabling direct connection and AI-powered penetration tests against Bitbucket-hosted repositories.

WAS THIS USEFUL? 👍 👎

runZero

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

runZero is a cloud-based asset discovery and management platform that identifies and inventories all connected devices across networks.

> runZero added a scan configuration toggle to enable or disable the promiscuous ARP probe.

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

01 **Promiscuous ARP probe toggle in scan config** NEW 45

A new configuration option on the scan configuration page lets users switch the promiscuous ARP probe on or off, useful for environments where broadcast-level ARP traffic must be minimized or is restricted.

Disable the promiscuous ARP probe for environments where broadcast traffic must be minimized.

> 📍 In the runZero console, go to the scan configuration page, locate the promiscuous ARP probe toggle, and switch it off before starting the scan.

Disable the promiscuous ARP probe for scans where broadcast-level ARP traffic is undesirable or restricted.

> 📍 In the runZero console, go to your scan configuration page, locate the Promiscuous ARP probe toggle, and switch it off before saving the scan.

— Names exact toggle and UI path, no deeper mechanism given

runZero 5.1.260825.0 (2026-08-25)

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Okta for AI Agents

Feedback ⓘ 1 RELEASE · 2026-08-26

CHANGELOG ↗

Okta for AI Agents provides identity and access management capabilities for AI agents, enabling secure authentication and authorization in AI-powered applications.

> Okta for AI Agents added Cross-App Access (XAA) so agents can act on resource apps over OIDC/SAML without runtime consent, replaced the Delegation tab with a new User access tab for configuring agent bindings, and shipped real-time import monitoring, new VPN IP categories, and a wave of new OIN integrations and API scopes.

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

01 **Cross-App Access (XAA) for AI agents** NEW 65

Admins can connect SSO-agentic requesting apps to SSO resource apps over OIDC and SAML without requiring user consent at runtime, with full admin visibility and control over which actions an AI agent can take on behalf of a user. Configured via the Admin Console under AI agents > select agent > User access > App used for access configuration, choosing Application as the resource type and selecting the XAA-enabled resource app.

Connect an AI agent to a resource app via XAA so the agent can act on behalf of users without runtime consent prompts.

> 📍 In the Admin Console, go to AI agents › select your agent › User access › App used for access configuration, choose Application as the resource type, then select your XAA-enabled resource app.

— Names OIDC/SAML mechanism and exact admin console navigation

changelog-20260826-ab6b5bf2

THINNER COVERAGE BELOW

02 **User access tab replaces Delegation tab for agent configuration**

55

BREAKING

The User access tab now configures AI agent access, replacing the previous Delegation tab delegation-link workflow; existing delegation links must be reconfigured using the User access tab. The same binding also lets an AI agent be bound to an OIDC SSO app

so they share the same credentials, with admins able to configure direct user authentication for the agent from that binding.

– Names both config tabs and binding mechanism, no reconfiguration steps

changelog-20260826-ab6b5bf2

03 **New API scopes for Harmony and Hero integrations** NEW

45

Adds `okta.schemas.read` scope to Harmony (API Service Integration), and `okta.roles.read`, `okta.userTypes.read`, `okta.logs.read`, and `okta.groups.read` scopes to Hero Integration.

– Names exact scopes; no usage guidance given changelog-20260826-ab6b5bf2

04 **New OIN catalog integrations added** NEW

40

New OIN integrations added: Airwallex (OIDC), Bold Group Stages (SAML), Gateco (SCIM), Orca Security (SAML), Orca Security (SCIM), Square (OIDC), Statsig Lifecycle Management Connector by Redblock (SCIM), Vimeo Lifecycle Management Connector by Redblock (SCIM), Briefly (OIDC), Briefly (SCIM), MicroWest Software Systems - AMMSWEB (SAML), MintMCP (SAML), Orchestra (OIDC), Orchestra (SCIM), Petual (SCIM), RansomLeak (SAML), RansomLeak (SCIM), and Remote.com (SCIM). Petual also gained Express Configuration support, and OCCAM Razor (OIDC) gained three new redirect URIs.

– Lists all named integrations but no setup steps changelog-20260826-ab6b5bf2

05 **Real-time progress in Import Monitoring dashboard** IMPROVED

30

The Import Monitoring dashboard now shows the number of data chunks currently being processed for in-progress imports.

– States what's shown but no interaction detail changelog-20260826-ab6b5bf2

06 **New IP service categories in enhanced dynamic zones** NEW

25

New IP service categories are now supported as individual VPN service categories in enhanced dynamic zones.

– Bare addition, no list of categories or config path changelog-20260826-ab6b5bf2

└──▷ **BREAKING ON UPGRADE**

! Orgs that previously used the Delegation tab to configure AI agent access through delegation links must reconfigure them using the User access tab.

WAS THIS USEFUL?  

ConductorOne

Feedback ⓘ 1 RELEASE · 2026-08-26 CHANGELOG ↗

ConductorOne is an identity governance platform that automates access management, provisioning, and compliance workflows across enterprise applications.

> ConductorOne shipped smart revocation controls that trace inherited access before approval, full push-rule execution history, new Copilot Studio and Gemini Enterprise integrations for the C1 MCP gateway, and c1i v0.5.0 with new subcommands, a breaking `--url` / `--server-url` flag split, and typed JSON output.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 `--server-url` flag replaces `--url` in c1i MCP commands

95

BREAKING

On `mcp servers register`, `mcp servers test-connection`, and `mcp servers update-credentials`, the external MCP server URL flag is renamed from `--url` to `--server-url`, so the global `--url` flag can now correctly select the C1 tenant instead. `c1i` now also requires a full `https://` tenant URL; the bare short-name shortcut is retired and a non-HTTPS URL is rejected instead of silently upgraded. Scripts or configs passing `--url` to these three commands will break.

Register a new external MCP server using the renamed `--server-url` flag and the global `--url` tenant selector.

```
$ c1i --url https://<tenant>.conductor.one mcp servers register --server-url https://my-mcp-server.example.com --name my-governed-tools
```

— Exact flag names, affected commands, and a runnable example are given.

changelog-20260826-b401642e

02 Push rule execution history per connector NEW

80

Push rules now record full execution history per connector, accessible via the Push rule page (all connectors) or View execution history in a push rule's row menu (single connector). Each entry shows who was pushed, when, which attributes changed, and the target system's rejection reason on failure.

— Exact UI navigation and entry fields named. changelog-20260826-b401642e

03 New c1i v0.5.0 subcommands for apps, auth, MCP gateway, and policies

NEW

78

c1i v0.5.0 adds `apps create`, `apps delete`, and `apps set-owners` subcommands rounding out app-management API coverage from the CLI; an `auth token` subcommand; `mcp gateway list-tools` and `mcp gateway call` subcommands; and a `policies` command family that guards against a silent deny-all default.

List all tools available through the C1 MCP gateway to verify governed tool exposure before wiring up a Copilot Studio or Gemini agent.

```
$ c1i --url https://<tenant>.conductor.one mcp gateway list-tools
```

— Command names given but full flags for each not shown. changelog-20260826-b401642e

04 **Smart revocation for inherited access** NEW

75

A new Revoke inherited access per-app setting enables smart revocation that traces access to upstream sources — direct grants, nested groups, role hierarchies, delegated provisioning — before an approver commits. It is off by default and shows a downstream-effects banner pre-confirmation.

— Mechanism and default state given but no API/config key named.

changelog-20260826-b401642e

05 **Configurable reviewer-visible attribute columns in access reviews**

NEW

75

Campaign admins can configure up to 10 reviewer-visible app user attribute columns (job title, department, cost center, etc.) in the access review table, available once the campaign's Default access review view is set to something other than None.

— Exact setting name and column limit given. changelog-20260826-b401642e

06 **New MCP gateway tool source integrations: Copilot Studio and Gemini Enterprise**

NEW

70

The C1 MCP gateway now supports connecting Copilot Studio agents, with every tool call authorized and logged under the calling user's identity and publishable to Microsoft 365 Copilot and Teams. It also supports Gemini Enterprise as a tool source, with every tool call attributed to the individual user rather than a shared service account.

List all tools available through the C1 MCP gateway to verify governed tool exposure before wiring up a Copilot Studio or Gemini agent.

```
$ c1i --url https://<tenant>.conductor.one mcp gateway list-tools
```

— Names both integrations and attribution behavior, no setup steps.

changelog-20260826-b401642e

---

07 **Public API endpoints for connector-discovered application management**

NEW 65

New public API endpoints list connector-discovered applications and their managed state, and promote an unmanaged application to managed — returning HTTP 409 on a repeat promotion, with the managed application ID readable from the binding.

— Behavior and status code given but exact endpoint paths not specified.

changelog-20260826-b401642e

---

08 **`Grant source` setting scopes automation revoke to direct grants**

NEW 65

Automation revoke steps gain a `Grant source` setting (value: `Direct`) to scope revocation to only direct grants, preventing collateral removal of group-inherited access.

— Names the exact setting and value but not full UI path. changelog-20260826-b401642e

---

09 **c1i list commands emit typed JSON and consistent exit codes**

IMPROVED 60

`c1i` v0.5.0 list commands now emit properly-typed Booleans, counts, and nullable fields instead of stringified or inconsistent JSON, and all responses map onto a single predictable exit-code scheme.

— Describes before/after but no exact schema or exit codes listed.

changelog-20260826-b401642e

THINNER COVERAGE BELOW

---

10 **C1AI assistant answers access history and grant-reason queries**

IMPROVED 55

The C1AI assistant can now answer access history queries — grants and revocations over a time period ordered by most recent — not just current access state, and returns the reason behind each grant (access request, group automation, catalog membership, uplift, etc.) alongside who has access.

— Describes capability but no example query or API. changelog-20260826-b401642e

11 **Filterable task list and revocation link on app overview**

IMPROVED

55

An app's task list tab can now be filtered to show Request tasks, and the Open revocations stat card on the app's Overview tab links directly to that filtered view.

— Names exact tabs and filter but limited scope. changelog-20260826-b401642e

12 **OCSF time field in automation circuit-breaker logs** IMPROVED

50

Automation circuit-breaker system-log events now include the standard OCSF time field, enabling direct SIEM ingestion without a parsing workaround.

— Names the field standard but gives no ingestion config example.

changelog-20260826-b401642e

13 **MCP servers searchable via global search** NEW

50

MCP servers are now searchable from global search (Cmd+K) by name, with recently opened servers appearing under Your recent searches.

— UI path given, minimal further detail. changelog-20260826-b401642e

14 **Read-only entitlement details drawer for reviewers** NEW

45

Any assigned reviewer can now open a read-only entitlement/resource details drawer from within a review task without requiring additional permissions.

— UI behavior named, no navigation detail. changelog-20260826-b401642e

15 **Past-grants search API returns addressable grant `id`**

IMPROVED

40

Each record returned by the public API's past-grants search now includes a unique `id` field, making individual grant history entries directly addressable.

— Field name given but no endpoint path or usage example. changelog-20260826-b401642e

16 **Toast notice for role-hidden page information** IMPROVED

35

Pages containing role-hidden information now display a toast — 'Some information on this page is hidden based on your permissions' — instead of showing a blank field with no indication.

— Exact message text given, minor UX change. changelog-20260826-b401642e

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! `c1i` now requires a full `https://` tenant URL; the bare short-name shortcut is retired and a non-HTTPS URL is rejected instead of silently upgraded.

! On `mcp servers register`, `mcp servers test-connection`, and `mcp servers update-credentials`, the external MCP server URL flag is renamed from `--url` to `--server-url`; scripts or configs passing `--url` to select the MCP server on these commands will break.

WAS THIS USEFUL?  

JumpCloud

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

Learn what password cracking is, how it works, and the best ways to protect against it.

> JumpCloud's August 25, 2026 agent release (2.159.1) adds a policy-controlled Self Service App Catalog for Windows and lifts the 64-character password limit on macOS, alongside smaller Device Admin and Device Trust refinements.

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

01 **Passwords longer than 64 characters on macOS** IMPROVED

53

Supports password changes longer than 64 characters on macOS, with a validation message shown on submission.

– Clear before/after limit but no exact new max or setting

JumpCloud Agent Release Notes for Augus…

02 **Self Service App Catalog UI via policy** NEW

52

Enables the Self Service App Catalog user interface on Windows via policy, giving managed users direct access to approved applications.

– Names UI and delivery method but no policy key or steps

JumpCloud Agent Release Notes for Augus…

03 **Fewer Device Trust certificate prompts on Windows** IMPROVED

35

Reduces Device Trust certificate selection prompts on Windows, speeding up device authentication.

– Names Device Trust and Windows but no mechanism detail

JumpCloud Agent Release Notes for Augus…

04 **Device Admin read-only state for single selection** IMPROVED

30

Device Admin controls now display a read-only state for single-selection cases, clarifying available actions.

– Thin UI note with no location or trigger detail

JumpCloud Agent Release Notes for Augus…

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 3 RELEASES · 2026-08-25 BLOG ↗

All the latest releases from the 1Password team

> 1Password's 8.12.34 releases add passwordless account linking across Linux, Mac, and Windows, graduate the MCP Server from Labs to general availability, and introduce an Offline Items sidebar view alongside new transparency and authentication options for 1Password for Claude.

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

01 MCP Server moves from Labs to Settings > Developer

BREAKING

80

The 1Password MCP Server has been moved out of `Settings > Labs` to `Settings > Developer` on Linux, Mac, and Windows, marking its graduation from experimental status to general availability.

Access the MCP Server configuration now that it has graduated out of Labs.

> 📍 In the 1Password app, go to Settings › Developer to find and configure the 1Password MCP Server.

Find and configure the 1Password MCP Server after it moved out of Labs.

> 📍 In the 1Password app, go to Settings › Developer and locate the MCP Server options (previously found under Settings › Labs).

— Names exact settings paths and gives navigation steps

1Password for Windows 8.12.34 (2026-08-… | 1Password for Mac 8.12.34 (2026-08-25) | 1Password for Linux 8.12.34 (2026-08-25)

02 **1Password for Claude access transparency and provider auth**

IMPROVED

65

Adds a button next to each item in the 1Password for Claude prompt that reveals why Claude is requesting access to that specific item, and enables 1Password for Claude to authenticate using logins configured with supported sign-in providers.

Review why Claude is accessing a specific credential before approving the request.

> 📍 In the 1Password for Claude prompt, select the button to the right of each item to view Claude's stated reason for requesting access.

— Includes UI step but omits provider list `1Password for Mac 8.12.34 (2026-08-25)`

THINNER COVERAGE BELOW

03 **Offline Items sidebar view** NEW

55

A new Offline Items entry in the sidebar surfaces items that cannot sync with a 1Password account, including deleted items, available on Mac and Windows and referenced for Linux.

Review items that failed to sync — including deleted ones — without leaving the main navigation.

> 📍 In the 1Password sidebar, select Offline Items to view all items that cannot be synced with your 1Password account.

— UI location given but sync mechanism unexplained `1Password for Mac 8.12.34 (2026-08-25)` `1Password for Windows 8.12.34 (2026-08-…` `1Password for Linux 8.12.34 (2026-08-25)`

04 **KDE Wayland auto-lock support** NEW

55

Adds auto-lock support in KDE Plasma on Wayland and other environments implementing the `ext_idle_notifier_v1` protocol.

— Names protocol but no configuration steps `1Password for Linux 8.12.34 (2026-08-25)`

05 **Passwordless account linking via 1Password.com** NEW

45

Enables adding a 1Password account via 1Password.com without re-entering a password, Secret Key, or SSO verification code, on Linux, Mac, and Windows clients.

— Explains behavior change but no setup steps given

`1Password for Mac 8.12.34 (2026-08-25)` `1Password for Linux 8.12.34 (2026-08-25)` `1Password for Windows 8.12.34 (2026-08-…`

06 **Clipboard history exclusion for account details** IMPROVED

Copying an account detail in the Manage Accounts pop-up no longer writes to the device clipboard history, reducing exposure of sensitive account data on Windows.

— Clear fix described but no reproduction steps `1Password for Windows 8.12.34 (2026-08-…`

└──▷ BREAKING ON UPGRADE

! The 1Password MCP Server setting has been removed from `Settings > Labs` and is now located under `Settings > Developer`.

WAS THIS USEFUL?  

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

Clerk is an authentication and user management platform that handles sign-up, login, and identity verification for applications.

> Clerk shipped two new capabilities: Backend API support for resetting a user's lost MFA factor, and an Admin Logs dashboard for auditing admin activity.

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

01 **MFA account recovery via Backend API** NEW 50

Adds MFA account recovery via the Backend API, enabling you to reset a user's second factor and build a custom MFA reset flow for users who have lost access to their authenticator.

— Names the Backend API but no specific endpoint or method given. product docs

02 **Admin Logs audit dashboard** NEW 40

New Admin Logs dashboard provides an audit trail of admin actions taken across a Clerk workspace, viewable in the Clerk Dashboard under Admin Logs.

Review recent admin activity across your workspace to investigate a suspicious change.

> 📍 In the Clerk Dashboard, go to Admin Logs to view a chronological audit trail of admin actions across your workspace.

— Gives a UI navigation path but no mechanism or scope detail.

Audit Dashboard activity with Admin Log…

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback

1 RELEASE · 2026-08-24

PRESS ↗

Okta Identity Engine is a modern identity platform providing authentication, authorization, and user management for applications.

> Okta made Agent SSO generally available, turning AI agents into first-class Universal Directory identities governed with short-lived tokens instead of static API keys, and expanded its broader Okta for AI Agents governance product alongside a wave of new Cross App Access integrations.

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

01 **Agent SSO general availability** NEW 85

Agent SSO is now generally available and included in core Okta SSO plans at no additional cost. It registers AI agents that support the Cross App Access protocol as first-class identities in Universal Directory alongside human employees, replacing static API keys and one-off OAuth grants with short-lived, identity-governed tokens. Administrators manage agent access policy through the same Okta Admin Console and workflows used for human employees — enabling assignment, monitoring, and policy updates for agents such as Anthropic Claude without separate tooling.

Govern an Anthropic Claude deployment's enterprise access through the Okta Admin Console without sharing static credentials or managing per-app OAuth grants.

📍 In the Okta Admin Console, go to Applications › Okta Integration Network, search for Anthropic Claude, add the integration, then navigate to the agent's profile in Universal Directory to assign access policies and monitor issued tokens.

— Names mechanism, plan inclusion, and gives exact admin console navigation.

launch-20260824-d920def3

02 **Okta for AI Agents governance and shadow AI discovery**

IMPROVED 80

Okta for AI Agents, a separate subscription generally available since May 2026, extends coverage beyond Cross App Access: it discovers shadow AI agents via browser, endpoint, and network detection; assigns named human owners; connects agents to custom authorization servers, service accounts, secrets, and other agents; and enforces governance through access certification reviews, approval workflows, and a kill-switch (agent deactivation). Agents onboarded through Agent SSO carry over automatically when upgrading to Okta for AI Agents, with no re-registration required.

— Detailed mechanism and scope, but no concrete API or config surface given.

launch-20260824-d920def3

03 **Cross App Access integrations in Okta Integration Network**

NEW 60

The Okta Integration Network now includes out-of-the-box Cross App Access integrations for Anthropic (Claude), `Archestra.AI`, Asana, Atlassian, Canva, Datadog, Figma, Glean, Granola, Linear, `MintMCP`, Notion, Slack, and Supabase, removing the need for custom integration work to onboard these agents/apps.

— Lists every named integration but no setup mechanism beyond that.

launch-20260824-d920def3

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

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

> DNSFilter added an exclusion operator to Query Log category filtering, several small dashboard usability improvements, and fixed a Relay Manager config-file installation issue on Windows.

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

01 **Dashboard visibility and navigation improvements** IMPROVED

57

Dashboard cards in the Manage Dashboards view now display Private or Shared visibility status, which can be updated directly from the quick actions menu; dashboards are now clickable directly from the grid view; and clicking the Total Requests widget now opens Insights with Allowed Requests already toggled on.

– Groups three thin UI tweaks in the dashboards area, each named

08/25 August 25 – Web Application and A…

02 **Relay Manager Windows config file location fix** IMPROVED

55

Relay Manager v0.2.2 defaults the config file installation to a shared system location for Windows deployments instead of the installing user's profile folder, preventing configuration loss when user profiles are removed.

– Explains before/after and reason, but no manual action given

08/25 August 25 – Web Application and A…

03 **'Does Not Include' operator in Query Log** NEW

53

Admins can now apply a Does Not Include operator to the Categories filter in the Query Log, letting them exclude specific categories from view during investigations.

– Names the exact filter and operator, no further mechanism given

08/25 August 25 – Web Application and A…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Varonis

Feedback ⓘ 9 RELEASES · 2026-08-03 → 2026-08-19 NOTES ↗

The leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis shipped a cluster of platform refinements this window: a new Data Classes model to simplify classification policy building, AWS Organization-scoped classification scans, faster Snowflake change detection and new sharing-lineage visibility, a unified Atlassian integration for Jira and Confluence, plus new coverage reporting, scan-pause scheduling, automation notifications, and connected-service-account filtering and export.

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

01 **Snowflake sharing lineage visibility** NEW 70

Varonis now displays the originating shares and associated accounts behind Snowflake resources labeled Shared Internally or Shared Externally, exposing the full inheritance chain behind each sharing relationship.

Investigate why a Snowflake resource is exposed externally by reviewing its originating shares and linked accounts directly in Varonis.

> 📍 In the Varonis dashboard, navigate to a Snowflake resource labeled Shared Externally, then open its sharing details to view the originating shares and associated accounts driving that classification.

— Names exact labels and includes navigation steps to investigate

See how and why Snowflake resources are…

02 **Data coverage reporting across scan scopes** NEW 65

A new built-in reporting view shows current and historical data coverage across configured scan scopes, broken down by cloud provider and service; trend analysis tracks coverage changes over time to validate scan scope configuration, and a coverage distribution view shows how monitored data is spread across cloud environments.

— Describes multiple report views but no exact navigation path

Understand data coverage across scan sc…

03 **Filtering and CSV export for connected service accounts** NEW 65

The Connected Services page now supports advanced criteria filtering for connected service accounts and export of the filtered results to CSV.

When auditing connected service accounts at scale, filter by advanced criteria and export the results to CSV for offline review or reporting.

> 📍 In the Varonis console, go to Connected Services › apply your advanced filter criteria › select Export to CSV.

— Exact page and export format with a runnable navigation sequence

`Filter and export connected service acc…`

THINNER COVERAGE BELOW

04 **Faster Snowflake masking policy and tag change detection**

IMPROVED

55

Snowflake masking policy and tag updates are now visible within 24 hours, replacing the previous scheduled-sync delay, so compliance teams can validate data protection controls more quickly.

— Concrete before/after and timing, but passive improvement with no user action

`Monitor Snowflake masking policy and ta…`

05 **Unified Atlassian integration for Jira and Confluence** NEW

55

A unified Atlassian integration connects both Jira and Confluence through a single onboarding flow, consolidating identities and eliminating duplicate users across Atlassian services.

— Names both integrated products but no setup steps

`Unify Jira and Confluence onboarding wi…`

06 **Data Classes for classification policies** NEW

50

Data Classes are a new way to build classification policies that group related patterns and dictionaries automatically instead of requiring individual selection, reducing policy complexity and false positives.

— Explains mechanism and benefit but no exact config path given

`Simplify classification policies with D…`

07 **AWS Organization scoping for classification scans** NEW

50

Scan scope rules now support AWS Organization ID, letting you include or exclude entire AWS Organizations without maintaining individual account inclusion/exclusion lists.

— Names the exact scoping capability but no UI steps

`Scope classification scans by AWS Organ…`

08 **Scan pause scheduling by day and time zone** NEW 50

Configurable per-day scan pause schedules with time zone selection let scan pauses align with maintenance windows and business hours across geographically distributed environments.

— Clear mechanism described, no exact configuration steps

Schedule scan pauses by day and time zo…

09 **Email alerts for pending automation actions** NEW 45

Varonis sends email notifications to customers when automation actions enter a pending state and require review in Action Center, enabling faster response to stalled workflows.

— Names Action Center but gives no configuration detail

Get notified when automation actions re…

WAS THIS USEFUL?  

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

RegScale is a compliance automation and GRC platform that helps organizations manage regulatory requirements, assessments, and security governance.

> RegScale's CLI 6.42.0 release adds new controls for Splunk evidence backfill and attribution, JCAM import mapping and dry-run reporting, and opt-in incremental data pulls from Wiz.

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

01 **Splunk evidence backfill and attribution settings** NEW

62

Adds a `splunkSids` setting for offline Splunk evidence backfill runs, along with Splunk evidence owner and approver settings for controlling evidence attribution.

— Names the config setting but no usage example given. CLI 6.42.0 (2026-08-26)

THINNER COVERAGE BELOW

02 **JCAM position-to-role mapping and dry run mode** NEW

58

Adds a JCAM position-to-role name mapping override to customize how positions are mapped during imports, plus a JCAM dry run mode that reports which positions' people would not be imported before committing the operation.

— Explains mechanism and purpose but no exact flag or command shown.

CLI 6.42.0 (2026-08-26)

03 **Opt-in Wiz incremental pulls** NEW

40

Adds opt-in Wiz incremental pulls that fetch only new or changed data instead of a full pull each time.

— Describes behavior but no config key or command for enabling it.

CLI 6.42.0 (2026-08-26)

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Giskard

Feedback ⓘ 2 RELEASES · 2026-08-26 NOTES ↗

Open-Source Evaluation & Testing library for LLM Agents

> Giskard shipped a ground-up v3 rewrite as a modular monorepo (`giskard-core`, `giskard-agents`, `giskard-checks`) that introduces a new Scenario/Suite testing API, a library of LLM-judge and validation checks, a beta OWASP red-teaming and RAG quality scanner, and a lean provider-agnostic LLM routing layer — alongside a patch letting `set_default_generator` accept model name strings directly.

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

01 **Scenario and Suite evaluation API** NEW 91

New `Scenario` API with .interact() and .check() chain lets you define multi-turn eval scenarios; `Suite` runs batches of scenarios with dynamic binding and supports Suite.append() chaining. Adds `annotations` support on scenario/trace objects for structured metadata, a print_report() method on result models with scenario and check error details in suite reports, JUnit XML export for `SuiteResult`, support for running individual scenarios multiple times, and pydantic-compatible input types for checks.

Define and run a grounded-answer eval scenario with a built-in LLM judge, then print the report.

```python
import asyncio
from giskard.checks import Scenario, Groundedness

def get_answer(inputs: str) -> str:
    return "Paris"

async def main() -> None:
    scenario = (
        Scenario("test_france_capital")
        .interact(inputs="What is the capital of France?",
outputs=get_answer)
        .check(
            Groundedness(
                name="answer is grounded",
                context="France is in Western Europe. Its capital
is Paris.",
            )
        )
    )
    result = await scenario.run()
    result.print_report()

asyncio.run(main())
```

— Full API surface named with a runnable example; no benchmarks. v3.0.0

02 **giskard-scan vulnerability and quality scanning** NEW

86

Introduces `giskard-scan` (beta) with `vulnerability_scan`, which red-teams an agent across OWASP LLM Top-10 threat categories including prompt injection, harmful content, stereotypes, and misinformation, and `quality_scan`, which evaluates RAG/knowledge-base quality. Installed via `pip install 'giskard[scan]'`.

Red-team an agent across OWASP LLM Top-10 categories including prompt injection and harmful content.

```python
import asyncio
from giskard.scan import vulnerability_scan

async def my_agent(inputs: str) -> str:
    return f"Echo: {inputs}"

async def main() -> None:
    await vulnerability_scan(
        target=my_agent,
        description="A customer support chatbot for an e-commerce platform.",
        languages=["en"],
    )

asyncio.run(main())
```

— Named functions, install command, categories, and runnable example. v3.0.0

03 **v3 monorepo rewrite and packaging** BREAKING 85

Rewrites Giskard as a v3 monorepo with packages `giskard-core`, `giskard-agents`, and `giskard-checks`; installs via `pip install giskard` (checks), `pip install 'giskard[scan]'` (vulnerability/quality scan), or `pip install 'giskard[openai]'` (provider SDK); requires Python 3.12+. Giskard v2 is no longer actively maintained, and the automatic tabular scan (`giskard.Model` + `giskard.Dataset`), `giskard.testing` ML test suite, and Giskard Hub are not present in v3.

— Exact package names, install commands, and removed v2 surfaces. v3.0.0

04 **Scenario and Interact API breaking changes** BREAKING 78

`Scenario.from_sequence` is removed, requiring step-based scenario construction via the .interact() chain. Jinja2 template parsing in `Workflow.chat` is now opt-in rather than default, templating support in conformity rules is dropped entirely, the `Interact` injection is now name-based (changing how check dependencies resolve inside scenarios), and the previous `Interact` API is replaced with a new API, breaking code that used the old interface.

— Enumerates every removed/changed surface with before/after. v3.0.0

05 **New LLM-judge and validation checks** NEW 75

`giskard-checks` adds four new checks: `AnswerRelevance` (LLM-judge check for whether agent responses are on-topic), `Toxicity` (LLM-judge check for harmful content), `RegexMatching` (pattern-based output validation with a regex timeout to mitigate ReDoS), and `JSONValidity` (asserts agent outputs are well-formed JSON).

Compose multiple checks with `AllOf` to require that agent output is both relevant and non-toxic.

```python
from giskard.checks import Scenario, AllOf, AnswerRelevance,
Toxicity

scenario = (
    Scenario("test_safe_response")
    .interact(inputs="Tell me about climate change.",
outputs=my_agent)
    .check(
        AllOf(
            AnswerRelevance(name="relevant"),
            Toxicity(name="non-toxic"),
        )
    )
)
```

— Names all four checks and their purposes; example shows two of them. v3.0.0

---

06 **Check composition operators AllOf, AnyOf, Not** NEW

70

`giskard-checks` adds `AllOf`, `AnyOf`, and `Not` check composition operators for building compound assertions, letting scenarios combine multiple checks into a single pass/fail condition.

Compose multiple checks with `AllOf` to require that agent output is both relevant and non-toxic.

```python
from giskard.checks import Scenario, AllOf, AnswerRelevance,
Toxicity

scenario = (
    Scenario("test_safe_response")
    .interact(inputs="Tell me about climate change.",
outputs=my_agent)
    .check(
        AllOf(
            AnswerRelevance(name="relevant"),
            Toxicity(name="non-toxic"),
        )
    )
)
```

— Named operators with usage example, but limited mechanism detail. v3.0.0

07 **Optional PostHog analytics with opt-out** NEW 70

Adds optional PostHog analytics to `giskard-core` and checks; opt out before importing via the `DO_NOT_TRACK=1` or `GISKARD_TELEMETRY_DISABLED=1` environment variables.

— Exact env vars given, immediately actionable, but no other detail. v3.0.0

08 **giskard-agents pipeline enhancements** NEW 67

`giskard-agents` gains a `metadata` parameter on the generator completion pipeline, a step-level type discriminator with tool input coercion and output serialization, generator retry and timeout policies (shared with `giskard-core`), and a generator-as-protocol-adapter pattern that lets generators serve as protocol bridges.

— Names several concrete surfaces but no usage examples given. v3.0.0

09 **set_default_generator accepts model strings** IMPROVED 60

`giskard-checks` v1.0.3 extends `set_default_generator` to accept model strings (e.g. a provider/model name) in addition to generator objects, simplifying default LLM generator configuration.

— Names the exact function and new accepted input type. giskard-checks/v1.0.3

THINNER COVERAGE BELOW

10 **giskard-llm provider-agnostic routing layer** NEW 50

Introduces `giskard-llm`, a lean provider-agnostic LLM routing layer that replaces litellm as the default LLM backend.

— States what it replaces but no usage or config detail. v3.0.0

11 **OWASP LLM suite generator** NEW 47

Adds a minimal OWASP LLM suite generator to `giskard-checks` covering LLM01 indirect injection.

— Named but described only in one thin line, no usage shown. v3.0.0

12 **UserSimulator personas and extended context** NEW 38

Introduces personas and extended context support for `UserSimulator` in the generators module.

— Bare feature name with no mechanism or example. v3.0.0

└──▷ **BREAKING ON UPGRADE**

! `Scenario.from_sequence` is removed; step-based scenario construction via the .interact() chain is now required.

! Jinja2 template parsing in `Workflow.chat` is now opt-in rather than the default, breaking any workflow that relied on automatic Jinja rendering of chat messages.

! Templating support in conformity rules is dropped entirely.

! The `Interact` injection is now name-based, changing how check dependencies are resolved inside scenarios.

! The `Interact` API is replaced with a new API (`refactor: new 'Interact' API`); code using the previous Interact interface will break.

! Giskard v2 is no longer actively maintained; the automatic tabular scan (`giskard.Model` + `giskard.Dataset`), `giskard.testing` ML test suite, and Giskard Hub are not present in v3.

WAS THIS USEFUL? 👍 👎

NVIDIA OpenShell

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

OpenShell is the safe, private runtime for autonomous AI agents.

> OpenShell v0.0.113 adds distributed tracing for Docker sandbox operations and corporate CA trust for Podman sandbox HTTPS proxies.

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

01 **OTLP trace export for Docker driver** NEW 40

Docker driver traces can now be exported over OTLP, enabling distributed tracing of sandbox container operations.

— Names OTLP and Docker driver but no config or flag given v0.0.113

02 **Corporate CA trust in Podman sandboxes** NEW 40

Podman-backed sandboxes now trust corporate CA certificates for HTTPS proxies and intercepted TLS.

— Names mechanism (CA trust, TLS interception) but no config key v0.0.113

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

Cloudflare is a web performance and security platform providing DDoS protection, CDN, DNS, and developer tools.

> Cloudflare shipped a broad update spanning AI, security operations and its REST API: new Workers AI models for video, image and transcription; Access service-token lifecycle controls and app-scoped API tokens; CASB auto-remediation and DLP test scanning; an RPKI ASPA path-validation tool on Radar; four WAF Managed Ruleset detections promoted from Log to Block (a breaking change) plus a new RCE rule; and dozens of new REST endpoints covering email security, Logpush transformers, API discovery, leaked credential checks, content scanning, AI Security for Apps, credential/scan management, Tunnel configuration and more.

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

01 **Radar RPKI ASPA path validation** NEW 90

Cloudflare Radar adds an ASPA validation tool in the Routing section: enter a BGP `AS_PATH` and check it against live RPKI ASPA records per `draft-ietf-sidrops-aspa-verification`, returning Valid/Invalid/Unknown verdicts for both the Upstream (section 5.4) and Downstream (section 5.5) algorithms, with a step-by-step verification trace and provider authorizations; path and algorithm are preserved in the URL for sharing. The `/bgp/rpki/aspa/snapshot` Radar API endpoint supplies the ASPA record snapshot used by the tool.

Fetch a live ASPA snapshot from Cloudflare Radar to power your own BGP route-leak analysis pipeline.

```
$ curl
https://api.cloudflare.com/client/v4/radar/bgp/rpki/aspa/snapshot
```

Validate a BGP AS path against live RPKI ASPA records to check whether a route announcement is authorized — useful when investigating suspicious routing events.

📍 In Cloudflare Radar, navigate to the RPKI ASPA path validation tool, enter the BGP AS path you want to verify, and review the verdict, provider authorizations, and step-by-step verification trace.

Validate a BGP AS path against live RPKI ASPA records on Cloudflare Radar to check route legitimacy during an incident.

> 📍 In Cloudflare Radar, navigate to the RPKI ASPA path validation tool, enter the BGP AS path to inspect, and review the verdict, provider authorizations, and step-by-step verification trace.

— Names the endpoint and spec sections, and gives a runnable curl example. product docs

---

02 **WAF Managed Ruleset promotes detections to Block** BREAKING

85

The Cloudflare Managed Ruleset promotes four detections from Log to Block: HTTP/2 Request Smuggling - Request Body Anomaly (rule `...1489d892`), XSS - JavaScript Event Handler Coercion - Headers (`...20646260`), - Body (`...d706d517`), and - URI (`...660886c8`) — meaning traffic previously only logged for these will now be blocked. It also adds a new Generic Rules - Remote Code Execution detection (`...2ca6cce3`) in Block mode, and merges the beta rule XSS, HTML Injection - Script Tag - Beta (`...c293b926`) into the original rule (`...7b58420b`), so overrides referencing the Beta rule ID no longer apply.

— Names every rule ID and the exact behavior change for each.

WAF - WAF Release - 2026-08-25 (2026-08…

---

03 **DLP Test scan capability** NEW

85

Cloudflare DLP adds a Test scan capability: paste text, upload a file, or upload a HAR file and select profiles to test directly against the DLP scanner, returning matched profiles, detection entries, confidence levels, match context, proximity keywords, file metadata, antivirus status, and OCR output — without routing traffic through Gateway or generating Gateway activity logs.

— Details full output fields and behavior, no exact navigation path given. product docs

---

04 **Access service token grace period and disable controls** NEW

80

Cloudflare Access lets administrators choose a grace period from 1 hour to 30 days when rotating a service token secret, keeping both the old and new secrets valid during rotation, with the API accepting an RFC 3339 expiration time for custom schedules; administrators can also temporarily disable a service token to halt authentication (including any active rotation grace period) without deleting its configuration, for later re-enablement.

— Explains rotation and disable mechanisms but no exact API endpoint named. product docs

---

05 **CASB auto-remediation policies** NEW

80

Cloudflare CASB now supports remediation policies that automatically revoke external sharing on Microsoft 365 and Google Workspace file-sharing findings, and/or send posture finding data to Slack, ServiceNow, or any webhook destination, without manual triage.

Set up a CASB remediation policy that auto-revokes external Google Workspace file shares and notifies a webhook the moment a finding is detected.

> 📍 In Cloudflare One, go to Cloud & SaaS findings › Policies › Create a policy. Under Basic information enter a Policy name. Under Choose how you want to trigger the policy select Vendor: Google, Integration: Google Workspace, and the desired Finding type. Under Define what to do with findings that match your trigger enable both Run Remediation and Send webhooks. Under Status enable Enable policy, then select Create policy.

— Explains mechanism and integrations, with exact UI steps to configure it. `product docs`

---

06 **Email security content policies and sending suppressions API**

NEW 80

Adds full CRUD (list, get, create, update, delete, batch) for email security content policies via `/accounts/{account_id}/email-security/settings/content_policies`, and full CRUD for account email sending suppressions via `/accounts/{account_id}/email/sending/suppressions`, including bulk import at `/accounts/{account_id}/email/sending/suppressions/bulk`.

Bulk-import email sending suppressions in one call to avoid sending to known-bad addresses across your account.

```
$ curl -X POST
'https://api.cloudflare.com/client/v4/accounts/{account_id}/email
/sending/suppressions/bulk' \
  -H 'Authorization: Bearer <API_TOKEN>' \
  -H 'Content-Type: application/json' \
  -d '{"suppressions": [{"email": "bounce@example.com", "reason":
"hard_bounce"}, {"email": "spam@example.com", "reason":
"spam_complaint"}]}'
```

— Names both endpoints with methods and gives a runnable bulk-import example.

`product docs`

---

07 **Logpush transformers API** NEW 80

Adds full CRUD API for Logpush transformers via `/accounts/{account_id}/logpush/transformers`, with additional endpoints to

preview a transformer (`/preview`), retrieve its content (`/{transformer_id}/content`), and list its versions (`/{transformer_id}/versions`).

Preview a new Logpush transformer to validate its output before saving it to a pipeline.

```
$ curl -X POST
'https://api.cloudflare.com/client/v4/accounts/{account_id}/logpu
sh/transformers/preview' \
  -H 'Authorization: Bearer <API_TOKEN>' \
  -H 'Content-Type: application/json' \
  -d '{"expression": "...", "sample": {}}'
```

— Names all four endpoint paths and includes a runnable preview call. product docs

---

08 **New Workers AI models: video, image and transcription** NEW

75

Adds `wan-3.0` , Alibaba's text-to-video model supporting adaptive aspect ratio and 480P/720P/1080P output with configurable duration; `universal-3.5-pro` , AssemblyAI's fast, high-accuracy speech recognition model; `flux-1-kontext-max` and `flux-1-kontext-pro` , Black Forest Labs' text-to-image and context-aware image editing models with strong character/style consistency; and `flux-video-upscale` , a video-to-video upscaler with precise and creative modes accepting clips up to 20 seconds with audio preserved.

— Names every new model ID and its parameters, no example call given. product docs

---

09 **Virtual Appliance downloads from the Connectors dashboard**

NEW

75

The Cloudflare One Virtual Appliance dashboard lets administrators select a hypervisor — VMware ESXi, Proxmox, or libvirt/KVM — and download the OVA image or install script directly from the Connectors page, removing the need to look up external asset URLs.

Download the Cloudflare One Virtual Appliance to deploy a connector in your on-premises hypervisor environment without hunting for an external download link.

> 📍 In the Cloudflare dashboard, go to Cloudflare One › Connectors, select your hypervisor type, then download the virtual appliance image or install script.

Download the Cloudflare One Virtual Appliance for your environment without leaving the dashboard.

> 📍 In the Cloudflare dashboard, go to Cloudflare One › Connectors, select your hypervisor, and download the virtual appliance image or install script.

— Gives exact dashboard path and supported hypervisors. product docs

10 **App-scoped API tokens for Flagship apps** NEW 70

Adds a Specified Flagship apps resource scope when creating custom API tokens, with Flagship App permission levels of Evaluate, Read, or Write, limiting a token's access to selected apps instead of the entire account.

— Names the exact scope and permission levels but not the API call. product docs

11 **MCP server portal supports stateless 2026-07-28 spec** IMPROVED 70

The MCP server portal `/mcp` endpoint now accepts stateless MCP `2026-07-28` requests alongside earlier 2025 Streamable HTTP clients, with automatic fallback negotiation for upstream servers; SSE connections continue using the legacy protocol.

— Names the endpoint and protocol version, but no migration steps. product docs

12 **New zone security settings endpoints** NEW 70

Adds Get NEL setting and Edit NEL setting endpoints at `/zones/{zone_id}/settings/nel`; endpoints to get and update session identifier settings; endpoints to get and update client-side security settings and to list, get, create, update, and delete content security rules; and endpoints to list and get detected connections, scripts, and cookies for client-side security.

— Names the NEL endpoint path; other settings endpoints lack explicit paths. product docs

13 **Selectable OAuth scopes in Wrangler login** IMPROVED 65

`wrangler login` and the Cloudflare API MCP server now present an OAuth consent dialog that lets users select individual optional scopes rather than approving all scopes; declining a scope and later needing it requires reauthorization.

— Names the exact command and behavior change but no scope list given. product docs

14 **Cloudflare Go SDK v7.9.0** IMPROVED 65

Bumps the Cloudflare Go SDK to `v7.9.0` (install via `go get -u github.com/cloudflare/cloudflare-go/v7@v7.9.0`) and the API spec to version `5.24.0`.

— Gives exact runnable install command and version numbers. product docs

THINNER COVERAGE BELOW

15 **API discovery operation labels and OpenAPI export** NEW

55

Adds API endpoints for API discovery operation labels — list, create, delete, get, update, and edit user-defined labels, plus managed label read and replace — along with endpoints to export discovered API operations as OpenAPI schemas and to list, edit, and manage discovered web and API operations.

— Describes operations added but no concrete paths named. product docs

16 **AI Security for Apps status and custom topics API** NEW

55

Adds Get AI Security for Apps status and Update AI Security for Apps status endpoints, plus endpoints to get and update custom topics for a zone.

— Names the status endpoints but not the custom-topics path. product docs

17 **API 403 responses include documentation_url** IMPROVED

50

Cloudflare API `403 Forbidden` responses now include a `documentation_url` field pointing directly to the API documentation for the denied endpoint.

— Names the exact field but nothing more to act on. product docs

18 **Token validation configuration and rules API** NEW 50

Adds token validation configuration endpoints to get, create, edit, and delete a configuration, replace or edit token validation credentials, and create, edit, get, or delete token validation rules.

— Lists operations but no endpoint paths given. product docs

19 **Leaked credential check API** NEW 50

Adds leaked credential check endpoints: get and update zone status, and list, create, get, update, and delete custom detection locations.

— Describes operations without concrete endpoint paths. product docs

20 **Content scanning API** NEW 50

Adds content scanning endpoints to enable, disable, and get or update status, and to list, create, and delete custom expressions for a zone.

— Describes operations without concrete endpoint paths. product docs

21 **Credential set and scan management API** NEW 40

Adds credential set and credential management endpoints (list, create, get, update, edit, delete) along with scan and target environment CRUD endpoints.

— Lists CRUD operations without naming any endpoint path. product docs

22 **WAF fallthrough expression template endpoint** NEW 35

Adds a Generate a fallthrough WAF expression template endpoint.

— Bare endpoint name, no path or usage shown. product docs

23 **Tunnel configuration API** NEW 35

Adds Get Tunnel configuration and Update Tunnel configuration endpoints.

— Names the endpoints but no path or usage detail. product docs

└─▷ BREAKING ON UPGRADE

! Rules `...1489d892`, `...20646260`, `...d706d517`, and `...660886c8` change from Log to Block — traffic previously only logged for HTTP/2 Request Smuggling and XSS JavaScript Event Handler Coercion will now be blocked.

! Rule `...c293b926` (XSS, HTML Injection - Script Tag - Beta) is merged into `...7b58420b`; any overrides or exceptions referencing the Beta rule ID will no longer apply.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 26, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue