THE CYBER TOOLCHAIN

NO. 079

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 19, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 105 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 13 tools

▾ // THE BRIEF HIDE

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

PROTECT

Cut the vulnerability queue by proving which findings can actually be
ched

vim ? help

AI Triage & Remediation Assist scores SAST/SCA findings for reachability and exploitability and generates candidate patches, running automatically on every scan or pull request. The manual work that goes away is an AppSec engineer reading each finding to decide whether the vulnerable path is live in this codebase before anyone writes a line of fix.

Checkmarx

DETECT

Run one alert queue across a fleet of serverless projects

Cross-project search lets an origin project query and investigate across up to 100 linked projects, so analysts triage, hunt and pivot in one place. Teams split across projects for tenancy or data-residency reasons previously had to context-switch per console and stitch the picture together by hand, which is how cross-environment activity gets missed.

Elastic Security Labs

GOVERN

Give every AI agent its own identity instead of a shared service account

Runlayer's Agent Accounts attach independent policies and an audit trail to each agent, so you can scope and revoke one agent without touching the rest and attribute an action to it after the fact. Exaforce's AI Security inventories AI apps and agents agentlessly with per-agent risk scoring and a kill switch — the prerequisite to both is knowing which agents exist and what each is allowed to do.

Runlayer · Exaforce

PROTECT

Publish to npm without a long-lived token that can push straight to the registry

Stage-only tokens and `npm stage publish` split staging a release from promoting it, so a leaked CI credential cannot silently ship a malicious version to consumers. This is the same failure mode behind most recent npm supply-chain incidents.

GitHub

APPLICATION SECURITY

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.165.13 NOTES ↗ RANK

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

> Infisical v0.165.13 introduces a project-wide unique-value validation rule, the first stage of recursive secret syncs, certificate filtering for PKI syncs, and a per-organization limit on secret scanning, alongside a new light mode.

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

01 **Unique-value validation rule for secrets** NEW 45

A new validation rule enforces that a secret value is not duplicated anywhere else within the project, checking uniqueness across the whole project rather than just within a single environment or folder.

– Explains the rule's scope but no config path given v0.165.13

02 **Certificate filters for PKI syncs** NEW 35

Adds the ability to filter which certificates are chosen when configuring PKI syncs.

– States the capability but not filter criteria or UI location v0.165.13

03 **Recursive secret syncs (stage one)** NEW 30

Secret sync jobs can now traverse nested folder structures recursively, described as the first stage of this capability.

– Named as stage one but mechanism and limits unspecified v0.165.13

04 **Per-organization secret scanning limit** IMPROVED 25

Secret scanning is now limited to one project per organization.

– Clear limit stated but no config or workaround detail v0.165.13

05 **Light mode** NEW 10

Adds a light mode to the interface, per the release summary; no further detail is provided.

– Mentioned only in summary with no supporting detail v0.165.13

v0.165.13 NOTES ↗ LATEST STABLE

> Infisical v0.165.13 adds light mode, recursive secret syncs, per-org secret scanning limits, PKI sync certificate filters, and a unique-values validation rule.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.165.13
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.165.13
```

Adds a 'unique values across the whole project' validation rule to enforce that a secret value is not duplicated in any other secret within the project.

Adds recursive secret syncs (stage one), enabling secret sync jobs to traverse nested folder structures.

Adds the ability to filter which certificates are chosen when configuring PKI syncs.

Limits secret scanning to one project per organization.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ

Copilot code review: An improved review experience (2026-09-18) BLOG ↗ RANK

Join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> GitHub's biggest shifts this window are a safer npm publishing model built around stage-only tokens and the `npm stage publish` command, and a broad refresh of Copilot code review's overview UI and automation, alongside a wide slate of Copilot agent features (model tiers, Dev Containers, Sentry canvas) and a new REST API for code coverage rulesets.

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

01 Stage-only npm tokens for gated publishing NEW 95

npm granular access tokens gain a Read and write (stage only) permission option, letting automated workflows stage package versions without direct-publish rights. The new `npm stage publish` command replaces `npm publish` for these workflows; staged versions require a maintainer to review and approve with 2FA before release, and npm rejects direct `npm publish` attempts even when bypass-2FA is configured. Stage-only tokens retain other write permissions such as moving dist-tags and deprecating versions. Staged publishing requires npm CLI 11.15.0+ and Node.js 22.14.0+.

Migrate a CI pipeline off bypass-2FA tokens by staging a release for maintainer approval instead of publishing directly.

```
$ npm stage publish
```

Create a stage-only granular access token scoped to your package so automation can submit releases without direct-publish rights — navigate to your npm account token settings and choose Read and write (stage only) for the target packages.

> 📍 In npm account settings, go to Access Tokens › Generate New Token › Granular Access Token, set permissions to Read and write (stage only), and scope it to the packages your workflow publishes.

— Names exact command, config option, and version requirements

Stage-only npm tokens for safer automat…

02 Copilot code review overview and automation improvements

IMPROVED

75

The Copilot code review overview comment now groups findings into Open, Resolved since last review, and Previously missed sections, each with severity and a link to the inline comment; every review comment gets a title reused in the overview issue list. Copilot generates a commit title and optional description automatically when committing an eligible batch of Copilot suggestions, including batches that mix in non-Copilot comments. Reviews also now resolve addressed comments during subsequent passes, suggest commit messages when applying suggestions, support shell tools to validate changes, and combine findings from multiple agents in Lite reviews.

▶ Watch video

Quickly triage a Copilot review by checking which issues are still open, which were fixed, and which were newly surfaced in existing code — all from the overview comment without scrolling through inline threads.

📍 In the pull request, open the Copilot overview comment › review findings grouped under Open, Resolved since last review, and Previously missed › click the inline-comment link next to any finding to jump directly to the relevant code.

— Names UI sections and behaviors but no API/flag

GitHub Copilot weekly releases — Septem… | Copilot code review: An improved review…

---

03 Efficiency, balance, and intelligence tiers for auto model selection

NEW

75

Auto model selection gains `efficiency`, `balance`, and `intelligence` tiers, letting users control how Copilot weighs cost, quality, and response time; rolling out in VS Code, Copilot CLI, and the Copilot app.

Prioritise response quality over cost for a high-stakes investigation by switching auto model selection to the `intelligence` tier.

📍 In VS Code (or Copilot CLI / Copilot app): open Copilot settings › Auto model selection › set tier to `intelligence` to have Copilot pick the highest-quality available model for each request.

— Names exact tier values and settings path GitHub Copilot weekly releases — Septem…

---

04 Code coverage ruleset condition via REST API NEW

70

The Restrict code coverage repository ruleset option, previously UI-only, is now exposed via the REST API, enabling programmatic create, update, and read of minimum line coverage percentage and maximum tolerable coverage drop rules for pull requests.

— Names the ruleset option but not exact endpoint path

Manage the code coverage ruleset condit…

05 **Sentry canvas in GitHub Copilot app** NEW 65

A Sentry canvas in the Copilot app lets users move from crash report to code fix: review errors, stack traces, and related context, then work with Copilot to investigate, validate a fix, and prepare a pull request.

When a Sentry alert fires in production, jump straight from the crash report into a fix workflow without leaving GitHub.

> 📍 In the GitHub Copilot app: open the Sentry canvas › select a Sentry error report › review the stack trace and related context › ask Copilot to investigate the root cause › validate the proposed fix › create a pull request from the canvas.

— Clear UI workflow steps but no config surface GitHub Copilot weekly releases — Septem…

06 **Local Dev Container support for agents in VS Code** NEW 65

Agents in the VS Code Agents window can now run with a project's own tools and dependencies via local Dev Container support, requiring Docker and a supported Dev Container configuration in the repo.

Run a Copilot agent against a project that requires specific runtimes or dependencies defined in a Dev Container, so the agent uses the same environment as local development.

> 📍 In VS Code 1.138, open the Agents window › start or resume a session › Copilot automatically launches the local Dev Container (Docker must be running and a Dev Container configuration must exist in the repo) › the agent executes tasks inside the container environment.

— Names requirements (Docker, Dev Container) and UI path

GitHub Copilot weekly releases — Septem…

07 **Budget request flow for Copilot usage-based billing** NEW 60

Copilot Business and Enterprise users on usage-based billing can request higher AI credit limits, which owners and billing managers can approve, adjust, or deny in settings; this is now generally available.

— Explains workflow and approvers but no exact UI path

GitHub Copilot weekly releases — Septem…

08 **Automatic inactive-session cleanup in VS Code Agents** NEW

60

An opt-in, preview feature in VS Code automatically marks Agents window sessions as Done once all their pull requests merge, with optional deletion after a configurable grace period.

— Describes mechanism and status but no config key named

GitHub Copilot weekly releases — Septem…

THINNER COVERAGE BELOW

09 **Pull request creation from Agent Host sessions** NEW

55

Pull requests can now be created directly from Agent Host sessions in the VS Code Agents window, including a generated title and description, draft status selection, and agent-assisted creation.

— Describes what and where but no example given GitHub Copilot weekly releases — Septem…

10 **VS Code Agents window usage metrics GA** IMPROVED

50

Usage metrics for the VS Code Agents window are now generally available, covering daily active users, session counts, user message totals, and user-level activity data in enterprise and organization reports.

— Lists metrics tracked but no access path given GitHub Copilot weekly releases — Septem…

11 **Copilot suggestions for repository custom property values**

NEW

40

Copilot can now suggest values for repository custom properties, in public preview, with organization and enterprise owners able to enable or disable the suggestions via a Copilot policy.

— Names the policy control but no exact key or path

GitHub Copilot weekly releases — Septem…

Copilot code review: An improved review experience (2026-09-18) BLOG ↗

LATEST STABLE

Copilot code review gains a structured findings overview, titled comments, smarter auto-resolution, and AI-generated commit messages for batch suggestions.

└──▷ **HOW TO FIND IT**

Quickly triage a Copilot review by checking which issues are still open, which were fixed, and which were newly surfaced in existing code — all from the overview comment without scrolling through inline threads.

📍 In the pull request, open the Copilot overview comment › review findings grouped under `Open`, `Resolved since last review`, and `Previously missed` › click the inline-comment link next to any finding to jump directly to the relevant code.

Adds a refreshed overview comment grouping findings into `Open`, `Resolved since last review`, and `Previously missed` sections, each with severity and a link to the inline comment.

Adds titles to every Copilot code review comment for at-a-glance prioritization, reused in the overview issue list.

Generates a commit title and optional description automatically when you commit an eligible batch of Copilot code review suggestions, including batches that mix in non-Copilot comments.

`Manage the code coverage ruleset condition with the REST API (2026-09-18)`

BLOG ↗ STABLE

GitHub REST API now supports managing the `Restrict code coverage` repository ruleset option programmatically.

`GitHub Copilot weekly releases — September 14 (2026-09-18)` **BLOG ↗** STABLE

Copilot gains Sentry canvas, Dev Container agents, efficiency/balance/intelligence model tiers, and GA usage metrics in VS Code.

└──▷ HOW TO FIND IT

When a Sentry alert fires in production, jump straight from the crash report into a fix workflow without leaving GitHub.

📍 In the GitHub Copilot app: open the Sentry canvas › select a Sentry error report › review the stack trace and related context › ask Copilot to investigate the root cause › validate the proposed fix › create a pull request from the canvas.

Run a Copilot agent against a project that requires specific runtimes or dependencies defined in a Dev Container, so the agent uses the same environment as local development.

📍 In VS Code 1.138, open the Agents window › start or resume a session › Copilot automatically launches the local Dev Container (Docker must be running and a Dev Container configuration must exist in the repo) › the agent executes tasks inside the container environment.

Prioritise response quality over cost for a high-stakes investigation by switching auto model selection to the `intelligence` tier.

📍 In VS Code (or Copilot CLI / Copilot app): open Copilot settings › Auto model selection › set tier to `intelligence` to have Copilot pick the highest-quality available model for each request.

Adds `efficiency`, `balance`, and `intelligence` tiers to auto model selection, letting users control how Copilot weighs cost, quality, and response time; rolling out in VS Code, Copilot CLI, and the Copilot app.

Adds a Sentry canvas in the GitHub Copilot app to move from crash report to code fix — review errors, stack traces, and related context, then work with Copilot to investigate, validate a fix, and prepare a pull request.

Adds local Dev Container support in the VS Code Agents window, letting agents run with a project's own tools and dependencies (requires Docker and a supported Dev Container configuration).

Adds pull request creation directly from Agent Host sessions in the VS Code Agents window, including generated title and description, draft status selection, and agent-assisted creation.

Adds Copilot suggestions for repository custom property values (now in public preview), with organization and enterprise owners able to enable or disable suggestions via a Copilot policy.

+4 more

Stage-only npm tokens for safer automation (2026-09-18) BLOG ↗ STABLE

npm granular access tokens now support `Read and write (stage only)` to gate package publishes behind 2FA reviewer approval.

└──▷ **TRY IT**

Migrate a CI pipeline off bypass-2FA tokens by staging a release for maintainer approval instead of publishing directly.

```
$ npm stage publish
```

Create a stage-only granular access token scoped to your package so automation can submit releases without direct-publish rights — navigate to your npm account token settings and choose `Read and write (stage only)` for the target packages.

📍 In npm account settings, go to Access Tokens › Generate New Token › Granular Access Token, set permissions to `Read and write (stage only)`, and scope it to the packages your workflow publishes.

> Adds Read and write (stage only) permission option when creating npm granular access tokens, allowing automated workflows to stage package versions without direct-publish rights.
>
> Introduces `npm stage publish` as the publishing command for stage-only token workflows, replacing `npm publish` — staged versions require a maintainer to review and approve with 2FA before release.
>
> Stage-only tokens retain other package write permissions (moving dist-tags, deprecating versions) but npm rejects direct `npm publish` attempts even when bypass-2FA is configured.
>
> Requires npm CLI 11.15.0 or later and Node.js 22.14.0 or later for staged publishing workflows.

WAS THIS USEFUL? 👍 👎

Checkmarx

Feedback ⓘ

Triage is where backlogs are born. Checkmarx Triage & Remediation Assist is how they're

BLOG ↗ RANK

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

> Checkmarx's headline addition this window is AI Triage & Remediation Assist, which evaluates SAST/SCA findings for reachability and exploitability, generates suggested code fixes, and can auto-triage on every scan or pull request; alongside it, Checkmarx shipped a new Software Supply Chain Security module (secrets detection, repository health scoring), GitHub App-based repo authentication, an MCP Server risk-triage surface, and a wave of DAST (proxy support, CLI workflow, authentication URL discovery, an Insights engine) and SCA/SBOM reporting improvements including license sections and CVE deduplication.

└──▷ **WHAT SHIPPED · 35 FEATURES** most completely described first ? what's the number?

01 **AI Triage and Remediation for SAST/SCA findings** NEW

100

A Triage with AI control in the Risk Orchestration side panel evaluates eligible SAST findings for reachability, exploitability, and existing sanitization, setting state to Confirmed or Proposed Not Exploitable with AI attribution; a Remediate with AI control generates a suggested code fix surfaced in the `Remediation > AI Remediation` tab with a what/why/how explanation, and can open a pull request for GitHub Code Repository Integration projects. Automated AI Triage can run on every completed scan, configurable by project, branch, scanner, risk status, and severity with administrator-controlled overrides, and also runs automatically on pull requests with direct remediation requests via PR comments for GitHub projects. Triage state changes propagate retroactively to identical SAST findings across earlier scans (matched by Similarity ID or Attack Vector ID) and forward into subsequent SCA scans (matched by package version and package manager, without retroactive application). An AI Triage & Remediation dashboard reports total triage and remediation actions, unique developers, estimated developer time saved, activity trends, triage outcomes, remediation results, and top projects, filterable by resource, branch, scanner, time range, tags, and groups. Both actions draw on Checkmarx Credits, with administrator visibility into burn rate and projected depletion; the AI Triage and AI Remediation agents were later upgraded for improved accuracy and performance, and new permissions were added for Analytics — AI Triage and Remediation.

Quickly assess whether a high-severity SAST finding is actually reachable and exploitable before assigning it to a developer.

> 📍 In Checkmarx One, open Risk Orchestration and select a SAST finding. In the side panel, choose `Triage with AI`. Review the reachability and exploitability analysis in the Info tab; the state is set to `Confirmed` or `Proposed Not Exploitable` with AI attribution visible.

After confirming a SAST risk, generate a review-ready code fix and open a pull request so the developer receives context-complete remediation guidance without leaving their normal workflow.

> 📍 In Checkmarx One, open the confirmed SAST finding in Risk Orchestration, select `Remediate with AI`, review the proposed change and what/why/how explanation in the `Remediation › AI Remediation` tab, then (for GitHub Code Repository Integration projects) trigger creation of a pull request containing the suggested fix.

Prevent backlog growth by automatically triaging all new Critical and High SAST findings on every main-branch scan, so analysts only review decisions rather than re-investigate context.

> 📍 In Checkmarx One administration settings, enable automated AI Triage and scope it to the target project, the main branch, the SAST scanner, and severity levels Critical and High. Set administrator-controlled overrides to restrict project teams from changing the branch and severity scope.

— Names exact controls, tabs, dashboard metrics, and gives runnable steps

`Triage is where backlogs are born.  Che…`

02 **License data in SCA Scan Reports** NEW 80

Revamped SCA Scan Reports add a `Package Licenses` section (CSV, XML, JSON) with per-package license data and a `Legal Risks` section (CSV, XML, JSON) with license usage risk data, available via the UI and REST APIs. Reports also gain license `Permissions`, `Limitations`, and `Conditions` data, plus an option to filter results to show only licenses marked as `Effective`.

— Names exact sections and export formats but not the REST endpoint path `product docs`

03 **Application-level SBOM reporting** NEW 75

Adds application-level SBOM report generation via the UI at `Workspace > Projects`, in addition to existing project-level SBOM support, with an option to exclude all licenses not designated as `Effective` for a package when generating SBOM reports.

— Names exact UI path and filter option `product docs`

04 **Secrets detection in Software Supply Chain Security** NEW

70

The Software Supply Chain Security (SCS) module adds a secrets scanner that identifies 170+ types of login credentials, access tokens, encryption keys, API keys, SSH keys, and webhook URLs across repositories. It was later extended with Custom Secret Detection Rules, letting organizations define, manage, and enforce tailored secret detection logic across codebases and workflows, and with Secret Detection added as a selectable scan type when generating reports.

— Names credential types and rules but no exact UI path `product docs`

---

05 **GitHub App-based authentication for repo integrations** NEW

70

Supports GitHub App-based authentication for code repository integrations, enabling short-lived scoped tokens, automatic token rotation, and granular permission control as a more secure alternative to Personal Access Tokens and OAuth apps.

— Explains token mechanism but no setup steps given `product docs`

---

06 **MCP Server natural-language risk triage** NEW 70

Enables changing a risk's state and severity directly through the MCP Server using natural language, without opening Checkmarx One — supports marking risks Not Exploitable or Confirmed across all major scan engines.

— Names exact states and surface but no sample command shown `product docs`

---

07 **Policy conditions for Container Security and SCA** NEW

65

Checkmarx One Policy Management gains new policy conditions per scanner: specialized Container Security conditions usable in Condition Groups, and a new `status` condition for the SCA scanner that triggers policies on New or Recurrent vulnerabilities while filtering out dev and test dependencies.

— Names the `status` condition and vuln states but not full syntax `product docs`

---

08 **Cloud Insights integrations for CrowdStrike and Sysdig** NEW

65

Cloud Insights can now connect to CrowdStrike's API endpoints via a setup wizard launched from the Cloud Insights page in the Checkmarx One UI, and Sysdig integration onboarding is streamlined by allowing account connection via a Spotlight token.

— Names UI entry point and token method but no full config steps `product docs`

---

09 **Full DAST workflow from the CLI** NEW 65

Enables running a full DAST workflow directly from the CLI via a REST-API-driven script that handles authentication, session creation, scanning, and results retrieval.

— Names the workflow stages but not the script or command itself `product docs`

10 **Attack Vector ID finding grouping** NEW 65

Introduces Attack Vector ID, a flow-aware finding grouping that encodes programming language, vulnerability query, and exact dataflow path — providing finer-grained separation than the default Similarity ID.

— Explains encoding and contrast with Similarity ID, no toggle location `product docs`

11 **Refined 'Not Calculated' result classifications** IMPROVED 60

Package results now distinguish Not Used (no usage detected) from Not Calculated (usage could not be determined), with a reason provided for Not Calculated entries (e.g., unsupported language). Risk results now similarly distinguish Not Found (no Exploitable Path detected) from Not Calculated (exploitability could not be determined), with a reason such as unsupported transitive dependencies.

— Details exact states and example reasons, no UI location given `product docs`

12 **DAST Insights engine for scan accuracy** NEW 60

Adds an Insights engine to DAST that flags conditions affecting scan accuracy — such as authentication failures, blocking mechanisms, or repeated server errors — as non-vulnerability findings; a High-level Insight stops the scan and marks it failed.

— Names conditions and severity behavior but not config to tune it `product docs`

13 **SCA dependency scanning: Direct scope and Java SBT** NEW 60

Adds SCA dependency scope tracking for Direct dependencies — packages called directly by the application — and adds Java SBT support (`build.sbt` , any `.sbt` file) for SCA dependency scanning.

— Names exact file type but no scan command or flag `product docs`

THINNER COVERAGE BELOW

14 **Repository health scoring in SCS** NEW 55

SCS adds a repository security health scanner that continuously evaluates repositories on code quality, dependency management, CI/CD best practices, and project maintenance. Repository Health was later added as a selectable scan type when generating reports, alongside Secret Detection, for a consolidated view across scan types.

— Describes evaluation dimensions but no scoring scale or thresholds `product docs`

15 **Azure AI guided remediation in IDE and CLI** NEW 55

Guided remediation in the Visual Studio Code Plugin and CLI can now integrate with a customer's Azure AI environment as a secure alternative to the OpenAI platform.

– Names surfaces but no setup steps or config keys product docs

16 **Status column for container image results** NEW 55

Adds a `Status` column to the results table indicating whether an image has been Muted, Snoozed, or Unresolved, with a tooltip explaining why an image wasn't resolved.

– Names exact states shown but not where results table is product docs

17 **ASPM risk score visibility in IDE and hover card** IMPROVED 50

Surfaces ASPM risk-scored results directly in the IDE, showing high-risk findings first to align with the ASPM model in the web app, and adds a card on hover over a vulnerability's risk score displaying the metrics, weighting factors, and contributing elements used to calculate it.

– Explains display mechanism but not exact IDE plugin or score formula product docs

18 **Proxy support for DAST** NEW 50

Adds proxy support for DAST, enabling scanning of internal, non-public, or firewall-protected applications from the cloud without whitelisting or exposing external IPs.

– States capability and benefit but not proxy config steps product docs

19 **Microsoft Teams webhook migration for feedback apps** BREAKING 50

Replaces deprecated Office 365 Connectors in the Microsoft Teams Feedback Apps integration with Microsoft Teams Workflows webhook URLs; customers must update their Teams webhook configuration after deployment or feedback notifications will stop being delivered.

– Clear migration requirement but no exact config field named product docs

20 **DAST authentication URL discovery** IMPROVED 50

DAST now prompts to include authentication URLs during setup and surfaces all discovered but non-included URLs in the authentication report, error messages, and configuration settings.

– Explains where URLs surface but not the exact setting name product docs

21 **TLS misconfiguration detection** NEW 50

Detects unsecured TLS connections and triggers an alert prompting configuration correction; findings appear in scan results and are automatically included in the default policy on environment setup.

— Explains detection and policy default but no severity or config key product docs

22 **Collapsible tabs in PR/MR decoration** IMPROVED 45

PR/MR decoration is extended with collapsible tabs across GitHub, GitLab, Bitbucket, and Azure DevOps (Bitbucket does not support collapsible tabs).

— Lists platforms and one limitation, no further mechanism product docs

23 **CVE deduplication for binary and source packages** IMPROVED 45

Deduplicates CVE results when a binary shares the same CVEs as its source package, counting them once as a unified result while still surfacing separate entries when binaries have distinct vulnerabilities.

— Explains dedup logic but no toggle or config to control it product docs

24 **SCA triage from the VS Code console** NEW 45

Enables triaging SCA results — editing state and adding comments — directly from the Visual Studio Code console.

— Names the surface but no exact commands product docs

25 **In-environment API collection management** NEW 45

Adds the ability to manage uploaded API collections and authentication scripts directly within an environment, available in both API and Web environments.

— States capability location but no UI path or command product docs

26 **Entities base report** NEW 40

Adds an Entities base report delivering insights into high-risk security issues and their trends over time, with flexible grouping and breakdowns by selected entity.

— Describes purpose but no report location or export format product docs

27 **Automated CVE publishing pending manual review** NEW 40

Adds an automated process that identifies and publishes newly identified CVEs with a pending manual review note, later updated after AppSec Research team analysis.

— Explains workflow but no location or trigger detail product docs

28 **Vulnerability distribution and local time zone in reports** IMPROVED 40

Reports gain a vulnerability distribution section with tables showing breakdown by status and state, improving governance reporting. Reports now also display date-time values in the user's local time zone (instead of UTC), with the applied time zone shown in the report header; browser time zone is applied automatically for UI-generated reports.

– Describes display changes with no export or config control named `product docs`

29 **Per-app ticket isolation in feedback apps** IMPROVED 35

Enhances feedback apps so each app only updates tickets it originally created, even when multiple apps share the same connection.

– Explains behavior but no config surface to control it `product docs`

30 **Bitbucket default branch auto-retrieval** IMPROVED 35

Automatically retrieves and applies the default branch when importing Bitbucket repositories by calling the Bitbucket API during import.

– Explains mechanism briefly but scope is narrow `product docs`

31 **Project scanning without repository integrations** NEW 30

Enables project scanning for projects that have no Code Repository integration configured.

– Bare description with no mechanism or setup detail `product docs`

32 **New 'inventory:attested' classification state** NEW 30

Adds `inventory:attested` state (`Inventory:attested`) as a recognized risk/inventory classification.

– Names the state but not how it's applied or triggered `product docs`

33 **SAST engine upgraded to 9.7.2** IMPROVED 25

Upgrades the SAST engine to version 9.7.2.

– Version number only, no changelog of engine changes `product docs`

34 **SCA CxScore composite risk score** NEW 25

Introduces a new SCA Composite Risk Score called CxScore.

– Name only, no calculation method or scale given `product docs`

35 **'Systemic' label for vulnerabilities** NEW 20

Adds a `Systemic` label for vulnerabilities.

— Bare name with no definition of what qualifies as systemic product docs

└──▷ **BREAKING ON UPGRADE**

! The Microsoft Teams Feedback Apps integration no longer uses Office 365 Connectors; customers must update to a Microsoft Teams Workflows webhook URL after deployment or feedback notifications will stop being delivered.

Triage is where backlogs are born. Checkmarx Triage & Remediation Assist is how they'

BLOG ↗ LATEST STABLE

Checkmarx One adds AI-powered Triage & Remediation Assist to evaluate reachability, exploitability, and generate suggested code fixes for SAST and SCA findings.

└──▷ **HOW TO FIND IT**

Quickly assess whether a high-severity SAST finding is actually reachable and exploitable before assigning it to a developer.

📍 In Checkmarx One, open Risk Orchestration and select a SAST finding. In the side panel, choose `Triage with AI`. Review the reachability and exploitability analysis in the Info tab; the state is set to `Confirmed` or `Proposed Not Exploitable` with AI attribution visible.

After confirming a SAST risk, generate a review-ready code fix and open a pull request so the developer receives context-complete remediation guidance without leaving their normal workflow.

📍 In Checkmarx One, open the confirmed SAST finding in Risk Orchestration, select `Remediate with AI`, review the proposed change and what/why/how explanation in the `Remediation › AI Remediation` tab, then (for GitHub Code Repository Integration projects) trigger creation of a pull request containing the suggested fix.

Prevent backlog growth by automatically triaging all new Critical and High SAST findings on every main-branch scan, so analysts only review decisions rather than re-investigate context.

📍 In Checkmarx One administration settings, enable automated AI Triage and scope it to the target project, the main branch, the SAST scanner, and severity levels Critical and High. Set administrator-controlled overrides to restrict project teams from changing the branch and severity scope.

Adds `Triage with AI` control in the Risk Orchestration side panel to evaluate eligible SAST findings for reachability, exploitability, and existing sanitization — setting state to `Confirmed` or `Proposed Not Exploitable` with AI attribution.

> Adds `Remediate with AI` control for eligible SAST findings, generating a suggested code fix surfaced in the `Remediation > AI Remediation` tab with a structured what/why/how explanation.
>
> Adds automated AI Triage on scan completion, configurable by project, branch, scanner, risk status, and severity, with administrator control over which parameters a project team can override.
>
> Adds an AI Triage & Remediation dashboard reporting total triage and remediation actions, unique developers, estimated developer time saved, activity trends, triage outcomes, remediation results, and top projects — filterable by resource, branch, scanner, time range, tags, and groups.
>
> Propagates AI Triage state changes to identical SAST findings retroactively across earlier scans (matched by Similarity ID or Attack Vector ID depending on account settings) and forward into subsequent SCA scans (matched by package version and package manager, without retroactive application).
>
> +1 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

Databricks Apps telemetry is now generally available (2026-09-18) BLOG ↗

RANK

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

> Databricks' biggest ship this window is Lakebase Search reaching GA with new Postgres extensions for vector and full-text search, alongside GA telemetry for Databricks Apps, beta support for base environments on classic compute, default Unity Gateway coverage for compliance security profiles, and beta account-only access to Genie One.

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

01 **Lakebase Search GA: vector and full-text extensions** NEW

98

Adds the `lakebase_vector` Postgres extension with the `lakebase_ann` index type for approximate nearest-neighbor vector search, drop-in compatible with pgvector, scaling to over 1 billion vectors on a single index with index builds up to 50-100x faster than HNSW. Adds the `lakebase_text` extension with the `lakebase_bm25` index type for BM25 full-text search over standard `tsvector` types, adding BM25 ranking and top-K pushdown not available in PostgreSQL's native GIN index. The two can be combined for hybrid search returning semantic and exact-term matches in a single ranked list, and indexes are storage-backed, available immediately after scale-to-zero with no warmup in serverless Lakebase projects.

Enable Lakebase Search in your project, then install both extensions to run hybrid semantic + keyword search against a single Postgres database.

```
$ -- After enabling Lakebase Search in project settings:
CREATE EXTENSION lakebase_vector;
CREATE EXTENSION lakebase_text;

-- Create a vector index (lakebase_ann) on an embeddings column
CREATE INDEX ON documents USING lakebase_ann (embedding);

-- Create a BM25 full-text index (lakebase_bm25) on a text column
CREATE INDEX ON documents USING lakebase_bm25 (body_tsvector);
```

Enable Lakebase Search through the project settings UI before installing the extensions.

> 📍 In your Lakebase project, go to Project Settings › Lakebase Search and enable the feature, then install `lakebase_vector` and/or `lakebase_text` via SQL.

— Runnable SQL (`CREATE EXTENSION`, index types) with concrete performance numbers

Lakebase Search is generally available …

THINNER COVERAGE BELOW

02 **Databricks Apps telemetry GA** NEW 50

Databricks Apps telemetry is now generally available, auto-persisting traces, logs, and metrics to Unity Catalog via OpenTelemetry. Once app telemetry is enabled, it supports custom instrumentation via the OpenTelemetry SDK alongside automatic system logs and usage event capture.

— Names OpenTelemetry SDK and Unity Catalog but no exact command

Databricks Apps telemetry is now genera…

03 **Account-only user access to Genie One (Beta)** NEW 50

Account-only users from a registered identity provider can now access account-level Genie One without being assigned to a workspace or granted consumer access, with automatic enrollment and Unity Catalog governance.

— Explains mechanism but no exact enablement steps

Account-only user access to Genie One (…

04 **Base environments on classic compute (Beta)** NEW 45

Base environments for managing Python dependencies, both Databricks-provided and custom workspace-level, are now supported on classic compute clusters in Beta.

— Describes the surface but no setup steps or config key

Base environments are now supported on …

> Databricks Apps telemetry is now generally available (2026-09-18) BLOG ↗ STABLE
>
> Databricks Apps telemetry now GA: traces, logs, and metrics auto-persisted to Unity Catalog via OpenTelemetry.
>
> Supports custom instrumentation via the OpenTelemetry SDK alongside automatic system logs and usage event capture once app telemetry is enabled.

> Lakebase Search is generally available (2026-09-18) BLOG ↗ LATEST STABLE

Lakebase Search GA adds `lakebase_vector` and `lakebase_text` Postgres extensions for scalable vector and BM25 full-text search.

└──▷ **TRY IT**

Enable Lakebase Search in your project, then install both extensions to run hybrid semantic + keyword search against a single Postgres database.

```
$ -- After enabling Lakebase Search in project settings:
CREATE EXTENSION lakebase_vector;
CREATE EXTENSION lakebase_text;

-- Create a vector index (lakebase_ann) on an embeddings column
CREATE INDEX ON documents USING lakebase_ann (embedding);

-- Create a BM25 full-text index (lakebase_bm25) on a text column
CREATE INDEX ON documents USING lakebase_bm25 (body_tsvector);
```

Enable Lakebase Search through the project settings UI before installing the extensions.

📍 In your Lakebase project, go to Project Settings › Lakebase Search and enable the feature, then install `lakebase_vector` and/or `lakebase_text` via SQL.

Adds the `lakebase_vector` Postgres extension with the `lakebase_ann` index type for approximate nearest-neighbor vector search, drop-in compatible with pgvector, scaling to over 1 billion vectors on a single index with index builds up to 50-100x faster than HNSW.

Adds the `lakebase_text` Postgres extension with the `lakebase_bm25` index type for BM25 full-text search using standard `tsvector` types and operators, adding BM25 ranking and top-K pushdown not available in PostgreSQL's native GIN index.

Supports hybrid search by combining `lakebase_vector` and `lakebase_text` to return semantic and exact-term matches in a single ranked list within one Lakebase database.

Delivers storage-backed indexes that are available immediately after scale-to-zero with no warmup required, enabling search in serverless Lakebase projects.

`Base environments are now supported on classic compute (Beta) (2026-09-17)`

BLOG ↗ STABLE

Databricks base environments for Python dependency management now supported on classic compute (Beta).

Supports Databricks-provided and custom workspace-level base environments for managing Python dependencies on classic compute clusters.

Unity Gateway support for compliance security profile standards (2026-09-17)

BLOG ↗ STABLE

Unity Gateway now available by default for workspaces with the compliance security profile enabled, covering all Databricks compliance standards.

Account-only user access to Genie One (Beta) (2026-09-16) BLOG ↗ STABLE

Genie One (Beta) now lets account-only users access account-level AI-BI without workspace assignment or consumer access grants.

WAS THIS USEFUL? 👍 👎

Elastic Security Labs

Feedback ⓘ

One SOC, 100 projects: running centralized alert triage on Elastic Security Serverless (1

BLOG ↗ RANK

Power insights and outcomes with The Elasticsearch Platform. See into your data and find answers that matter with enterprise solutions designed to help you accelerate time to insight.

> Elastic Security Serverless introduced cross-project search (CPS), letting one origin project centralize alert triage, hunting, and investigation across up to 100 linked projects.

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

01 Cross-project scope in Discover and dashboards NEW

83

Discover and non-managed dashboards can now run ES|QL queries and rollups across all linked projects for broad hunting sweeps, using an unqualified `FROM` that automatically includes newly linked projects with no query edits. On Security managed pages, a CPS scope selector shows in-scope projects as read-only, while admins configure default scope per space.

Hunt for an entity's activity across every linked project at once using an unqualified FROM in Discover — new projects linked later are automatically in scope with no query edits.

esql

```
FROM logs-* METADATA _index
| WHERE user.name == "<username>"
| KEEP _index, @timestamp, event.action, source.ip, user.name
| SORT @timestamp DESC
```

Build situational-awareness dashboards by rolling up alert counts per linked project and severity — validated interactive across 100 linked projects.

esql

```
FROM .alerts-security.alerts-* METADATA _index
| STATS count = COUNT(*) BY _index, kibana.alert.severity
| KEEP _index, kibana.alert.severity, count
| SORT count DESC
```

— Includes runnable ES|QL examples and exact scoping behavior

One SOC, 100 projects: running centrali…

02 **Cross-project search for centralized alert triage** NEW

65

A new capability lets one origin project run its roughly 2,100 detection rules once and have alerts from up to 100 linked projects consolidated into a single queue on the origin, eliminating per-project rule duplication and the need for analysts to hop between project contexts.

— Explains mechanism and scale but no setup steps given

One SOC, 100 projects: running centrali…

03 **Cross-project Timeline, Alert flyout, and Analyzer views** NEW

60

Timeline search, the Alert flyout, and the Analyzer/Correlations views now extend (read-only) across linked projects, so analysts can follow an attack chain — such as failed console logins followed by `AttachUserPolicy` and `CreateAccessKey` — without leaving the origin project.

— Names specific views and an example chain, no config steps

One SOC, 100 projects: running centrali…

THINNER COVERAGE BELOW

04 **Machine learning datafeeds across linked projects** NEW

50

ML datafeeds can now read data from linked projects, with jobs and their results stored on the origin project, enabling centralized ML anomaly detection across the full set of linked projects.

— Explains mechanism but omits configuration detail

One SOC, 100 projects: running centrali…

05 **Entity store partial cross-project support** NEW 45

Origin-project entity profiles now include entities from every linked project, but per-project risk scores are not merged across projects — a stated limitation of this initial cross-project entity store support.

— States the limitation clearly but no workaround

One SOC, 100 projects: running centrali…

06 **Default limit of 100 linked projects** NEW 40

Cross-project search enforces a default limit of 100 linked projects per origin project; higher limits are available by contacting Elastic.

— Gives exact limit and path to raise it One SOC, 100 projects: running centrali…

One SOC, 100 projects: running centralized alert triage on Elastic Security Serverless

BLOG ↗ LATEST STABLE

Elastic Security Serverless gains cross-project search (CPS), letting one origin project run ~2,100 detection rules across up to 100 linked projects.

└──▷ **USE IT**

Hunt for an entity's activity across every linked project at once using an unqualified FROM in Discover — new projects linked later are automatically in scope with no query edits.

esql

```
FROM logs-* METADATA _index
| WHERE user.name == "<username>"
| KEEP _index, @timestamp, event.action, source.ip, user.name
| SORT @timestamp DESC
```

Build situational-awareness dashboards by rolling up alert counts per linked project and severity — validated interactive across 100 linked projects.

esql

```
FROM .alerts-security.alerts-* METADATA _index
| STATS count = COUNT(*) BY _index, kibana.alert.severity
| KEEP _index, kibana.alert.severity, count
| SORT count DESC
```

Consolidates all rule-generated alerts into a single queue on the origin project, eliminating per-project rule duplication and multi-context analyst hopping.

Surfaces cross-project scope in Discover and non-managed dashboards for broad ES|QL hunting sweeps; Security managed pages show the CPS scope selector as read-only so analysts see in-scope projects while admins set defaults per space.

Extends Timeline search, Alert flyout, and Analyzer/Correlations views (read-only) across linked projects so analysts can follow an attack chain without leaving the origin.

Supports machine learning datafeeds reading linked-project data, with jobs and results stored on the origin for centralized ML anomaly detection.

Provides entity store partial cross-project support: origin profiles include entities from every linked project, though per-project risk scores are not merged across projects.

+1 more

WAS THIS USEFUL?  

Exaforce

Feedback ⓘ announced · Sep 15, 2026 PRESS ↗ RANK

Exaforce runs AI Exabots that detect, triage, investigate and respond to threats across cloud, SaaS, identity, endpoint and code on a real-time knowledge graph.

> Exaforce launched AI Security, expanding its Exabot platform beyond the SOC with agentless inventory of AI apps and agents, per-agent risk scoring, governance controls, correlated threat detection, and containment actions including a kill switch for rogue agents.

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

01 **Threat containment actions including agent kill switch** NEW

80

Adds threat containment actions — session revocation, model-provider API key deactivation, device isolation, and an agent kill switch (ending an agent's process) — executed through existing EDR, identity, and model-provider controls, with autonomy tunable per action from analyst-approved to fully autonomous.

— Names concrete containment actions and tunable autonomy but no exact API/CLI.

launch-20260915-8c1d4374

02 **Agentless inventory of AI apps and agents** NEW 65

Adds a continuous, agentless inventory of AI apps connected to company accounts, coding agents (Claude Code , Cursor), hosted agents (custom GPTs), MCP servers, skills, and IDE plugins, tied to the people and permissions behind them — with no new endpoint agent, browser extension, or gateway to deploy.

— Names specific tools covered but no configuration surface to act on.

launch-20260915-8c1d4374

03 **Correlated AI threat detection across identity, endpoint and code**

NEW 60

Adds threat detection that correlates agent actions, agentic-app activity, and prompts with human identities, endpoint data, file access, and code context to detect misuse, sensitive-data exposure, and active threats — including chains of individually legitimate-looking actions.

— Describes detection mechanism but gives no rule names or thresholds.

launch-20260915-8c1d4374

04 **Per-agent and per-skill risk scoring** NEW 60

Adds per-agent and per-skill risk scoring for every agent, skill, MCP server, and plugin, paired with a recommended fix, covering exposures such as excessive permissions and unvetted MCP servers.

— Names exposure types and recommended-fix mechanism, no scoring scale given.

launch-20260915-8c1d4374

THINNER COVERAGE BELOW

05 **Expanded knowledge graph data ingestion** NEW 55

Ingests EDR endpoint data, model-provider audit and usage logs, and productivity-suite activity into the Exaforce knowledge graph alongside existing identity, cloud, SaaS, and code data.

— Lists data sources ingested but no schema or query detail. launch-20260915-8c1d4374

06 **Cross-provider AI model coverage** IMPROVED 50

Adds support for OpenAI, Gemini, and Microsoft Copilot model providers to the existing Claude Compliance API integration, giving teams cross-provider AI activity coverage on one platform.

— Names providers and existing API integration but no setup steps.

launch-20260915-8c1d4374

07 **AI app and agent governance controls** NEW 45

Adds governance controls to define which AI apps, agents, and models are sanctioned and for whom, enforced through existing identity, endpoint, and model-provider controls.

— States purpose and enforcement path but no policy syntax or UI location.

launch-20260915-8c1d4374

08 **Token-usage anomaly monitoring** NEW 40

Adds token-usage monitoring with anomaly detection to identify runaway AI token use, exploitation, or misuse across the organization.

— Describes purpose only, no thresholds or alerting mechanism given.

launch-20260915-8c1d4374

09 **Personal vs business AI usage classification** NEW 30

Adds visibility into chats and prompts classified as personal or business use to support adherence to usage policies.

— Brief description with no classification method or policy detail.

launch-20260915-8c1d4374

announced · Sep 15, 2026 PRESS ↗

Exaforce AI Security adds agentless visibility, risk scoring, governance, threat detection, and an agent kill switch for AI agents and agentic apps.

Adds support for OpenAI, Gemini, and Microsoft Copilot model providers to the existing Claude Compliance API integration, giving teams cross-provider AI activity coverage on one platform.

Adds a continuous, agentless inventory of AI apps connected to company accounts, coding agents (Claude Code, Cursor), hosted agents (custom GPTs), MCP servers, skills, and IDE plugins, tied to the people and permissions behind them — with no new endpoint agent, browser extension, or gateway to deploy.

Adds visibility into chats and prompts classified as personal or business use to support adherence to usage policies.

Adds per-agent and per-skill risk scoring for every agent, skill, MCP server, and plugin, paired with a recommended fix, covering exposures such as excessive permissions and unvetted MCP servers.

Adds governance controls to define which AI apps, agents, and models are sanctioned and for whom, enforced through existing identity, endpoint, and model-provider controls.

+4 more

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ Introducing MSP Billing NOTES ↗ RANK

Manage devices and everything they touch. See how security, IT, and compliance get easier when unified around devices with Iru.

> Iru added MSP Billing for partners in its Partner Portal and introduced per-device Windows patch compliance status in Windows Managed OS Library Items.

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

01 **Windows patch compliance status per device** NEW 70

Windows Managed OS Library Items now report per-device patch compliance with Pass, Update past due, or Error status. Expanding a device row reveals its current Windows 11 build versus the required build, plus eligibility, install enforcement, and restart enforcement dates.



— Describes exact statuses and data shown, with clear UI navigation.

Windows Managed OS patch compliance

02 **MSP Billing in Partner Portal** NEW 60

Adds MSP Billing in the Partner Portal, giving U.S. `-based` managed service providers daily, per-tenant spend visibility and the ability to update payment methods, billing information, and view past invoices.

.

— Names concrete capabilities and UI location but no config/API surface.

Introducing MSP Billing

Introducing MSP Billing NOTES ↗ STABLE

Iru adds MSP Billing for managed service providers with per-tenant spend visibility and self-serve setup in the Partner Portal.

Windows Managed OS patch compliance NOTES ↗ LATEST STABLE

Iru Windows Managed OS Library Items now report per-device patch compliance with Pass, Update past due, or Error status.

Expanding a device row reveals its current Windows 11 build versus the required build, plus eligibility, install enforcement, and restart enforcement dates.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

ConductorOne

Feedback ⓘ September 18, 2026 (2026-09-18) BLOG ↗

RANK

Govern workforce access, deploy AI agents, and get work done with C1.

> ConductorOne's biggest changes this window are on the API surface: a public REST API for AI governance, new Terraform resources for MCP sources, and C1AI task diagnostics, alongside several breaking changes that split secret deletion from credential revocation and remove the spend-remedy status endpoint.

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

01 **Secret deletion and credential revocation split** BREAKING 95

`DELETE /api/v1/secrets/{vault_id}` now deletes only the secret and no longer revokes its credential; use the new `POST /api/v1/app-secrets-admin/revoke` endpoint to queue a credential revoke separately from secret deletion. Service SDKs rename `Revoke` to `Delete` to match.

— Exact endpoints and required migration step named. September 18, 2026 (2026-09-18)

02 **Spend-remedy task API removed** BREAKING 85

`GET /api/v1/tasks/{task_id}/spend-remedy-status` has been removed — migrate to the task APIs. JS/TS clients must handle spend-control revision values as strings, and integrations with a closed list of task action types must handle the new `TYPE_SPEND_REMEDY = 7`.

— Exact removed endpoint and required client changes. September 18, 2026 (2026-09-18)

03 **C1AI task diagnosis for stuck tasks** NEW 85

Super Admins can now ask C1AI to diagnose a specific stuck, failed, or waiting task; the assistant traces the task's approval history, provisioning logs, access-profile origin, and connector handoff, then recommends whether to approve, deny, or retry it. For example, asking 'why is task #4821 stuck?' returns the approval and provisioning state, access profile origin, app-account access, and connector handoff details.

When a provisioning task is stuck or shows unexpected access state, ask C1AI to trace its full history and get a recommended remediation step before manually acting.

📍 In the C1AI assistant, type: 'why is task #4821 stuck?' — the assistant returns its approval and provisioning state, access profile origin, app-account access, and connector handoff details, then recommends whether to approve, deny, or retry the task.

— Includes a runnable example query and full diagnosis trace.

September 18, 2026 (2026-09-18)

04 **Resource-type filters and revoke-criteria matching** BREAKING 75

Resource-type filters in Grant Found, Grant Deleted, and revoke criteria now use a faceted, searchable picker with real counts and name disambiguation. A revoke criteria scoped only by a legacy resource-type name now fails the step instead of silently matching every app sharing that name — re-scope any automation relying on name-only matching.

— Names affected steps and required remediation. September 18, 2026 (2026-09-18)

05 **Findings rules automation and assignment** NEW 70

Findings rules can now automatically set severity, add or remove annotations, or assign every matching finding to a specific person as it arrives, via a new Set assignee step in Transformation rules; a person's own manual assignment always overrides a later rule.

— Names the rule step but not full config path. September 18, 2026 (2026-09-18)

06 **Public REST API for AI governance** NEW 60

The public REST API now exposes classifiers, tool gates, classifier bindings, templates, and the agent classifier policy for managing AI governance programmatically.

— Names governed objects but no endpoint paths given. September 18, 2026 (2026-09-18)

THINNER COVERAGE BELOW

07 **My resources page consolidation** IMPROVED 55

The My resources page gains Apps and Entitlements tabs, consolidating groups, vaults, credential stores, secrets, apps, and entitlements — including secrets shared with you — into one view.

— Names the page and tabs but no further mechanism. September 18, 2026 (2026-09-18)

08 **Terraform resources for MCP sources** NEW 55

New Terraform resources cover hosted and external MCP sources, including their connection details, authentication, and tool enablement policy.

— Names resource scope but not exact resource types. September 18, 2026 (2026-09-18)

09 **Guided setup flows for apps and MCP integrations** IMPROVED

55

A single guided setup flow now replaces the multi-page form for creating apps and adding connectors. Registered MCP integrations (catalog and external) also share one guided setup flow covering connection, access, tool approval, and authentication.

— Names both flows' scope but no exact steps. September 18, 2026 (2026-09-18)

10 **Connector config change justification** NEW

50

Admins can now require a justification whenever a connector's configuration changes, and changes are tracked in a secret-safe configuration history.

— Describes the control but not where to enable it. September 18, 2026 (2026-09-18)

11 **MCP access requires explicit toolset grant** BREAKING

45

MCP access now requires an explicit grant on a server's toolset — app membership or admin role alone no longer implies access.

— States the new rule without migration steps. September 18, 2026 (2026-09-18)

12 **Denial justification in access-review revokes** IMPROVED

45

Access-review revoke tasks now automatically include the reviewer's denial justification, surfaced in ticketing-system templates via new template variables.

— Names template variables generically, not the exact names.

September 18, 2026 (2026-09-18)

13 **c1aw agent discovery over C1 MCP endpoint** NEW

40

Customer agents connected to your C1 MCP endpoint can now discover and converse with c1aw agents, with admin-managed sessions.

— States the capability without setup detail. September 18, 2026 (2026-09-18)

14 **Automation step and trigger search** NEW

40

Automation builders can now search available steps and trigger types by name or description.

— Simple UI search capability, minimal detail. September 18, 2026 (2026-09-18)

15 **Conflict monitor group size handling** IMPROVED 35

Conflict monitors can now handle groups of very different sizes, as long as the total stays under 1,024 members.

— Gives the limit but no other mechanism. September 18, 2026 (2026-09-18)

16 **Campaigns scope pending invitations** IMPROVED 35

Campaigns can now include users with pending (unaccepted) account invitations in their scope.

— Simple scope change with no further detail. September 18, 2026 (2026-09-18)

17 **Resource description column in review tables** NEW 35

Reviewers can now see a `Resource description` column in review tables, showing custom descriptions or connector-synced metadata.

— Names the UI column but nothing further. September 18, 2026 (2026-09-18)

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! A revoke criteria scoped only by a legacy resource-type name now fails the step instead of silently matching every app sharing that name — re-scope any automation relying on name-only matching.

! `DELETE /api/v1/secrets/{vault_id}` now deletes only the secret and no longer revokes its credential; use the new `POST /api/v1/app-secrets-admin/revoke` to queue a credential revoke separately. Service SDKs rename `Revoke` to `Delete` to match.

! `GET /api/v1/tasks/{task_id}/spend-remedy-status` has been removed — migrate to the task APIs. JS/TS clients must handle spend-control revision values as strings, and integrations with a closed list of task action types must handle the new `TYPE_SPEND_REMEDY = 7`.

September 18, 2026 (2026-09-18) BLOG ↗ LATEST STABLE

ConductorOne adds C1AI task diagnostics, finding assignment, MCP Terraform resources, and a public REST API for AI governance

Adds `POST /api/v1/app-secrets-admin/revoke` endpoint to queue a credential revoke separately from secret deletion.

Exposes classifiers, tool gates, classifier bindings, templates, and the agent classifier policy through the public REST API for AI governance management.

Adds Terraform resources for hosted and external MCP sources, covering details, authentication, and tool enablement policy.

Adds a `Set assignee` step in Transformation rules to automatically assign a finding to a specific person.

Super Admins can now ask C1AI to diagnose a specific stuck, failed, or waiting task, tracing approval history, provisioning logs, access-profile origin, and connector handoff.

`+12 more`

└─▷ **BREAKING ON UPGRADE**

- ! A revoke criteria scoped only by a legacy resource-type name now fails the step instead of silently matching every app sharing that name — re-scope any automation relying on name-only matching.
- ! `DELETE /api/v1/secrets/{vault_id}` now deletes only the secret and no longer revokes its credential; use the new `POST /api/v1/app-secrets-admin/revoke` to queue a credential revoke separately. Service SDKs rename `Revoke` to `Delete` to match.
- ! `GET /api/v1/tasks/{task_id}/spend-remedy-status` has been removed — migrate to the task APIs. JS/TS clients must handle spend-control revision values as strings, and integrations with a closed list of task action types must handle the new `TYPE_SPEND_REMEDY = 7`.

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ Actions - Modules - GA (2026-09-18) BLOG ↗ RANK

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

> Auth0 shipped general availability of Actions Modules, letting developers share reusable code, secrets, and dependencies across Actions and Triggers, alongside a new npm package for testing and editor support.

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

01 **Actions Modules reach general availability** NEW 85

Auth0 Actions Modules is now GA, letting you create, manage, and share reusable code across multiple Actions and Triggers within an Auth0 Tenant. Modules support independent secrets and dependencies decoupled from the Actions that import them, enable cross-trigger reuse of custom functions by importing a single Module into any Action regardless of which Trigger it is bound to, and improve execution performance by moving expensive initialization logic into a Module that is initialized once and reused across Action invocations.

Create a new Action Module in the Auth0 Dashboard to house shared initialization logic (e.g. an SDK client) that multiple Actions can import instead of each re-running setup code.

> 📍 In the Auth0 Dashboard, go to Actions › Modules › Create Module, give it a name, add any required secrets and npm dependencies, then write and publish your shared code.

— Describes mechanism and gives dashboard steps, but no API/CLI syntax

Actions - Modules - GA (2026-09-18)

THINNER COVERAGE BELOW

02 **@auth0/actions npm package for testing** NEW 57

Adds the `@auth0/actions` npm package providing unit-testing helpers and IntelliSense (autocompletion and type hints) for Actions Modules in any editor.

— Names the package but gives no install or usage command

Actions - Modules - GA (2026-09-18)

└──▷ **ALSO FROM THESE RELEASES**

. . . .

> Actions - Modules - GA (2026-09-18) BLOG ↗ LATEST STABLE

Auth0 Actions Modules reaches GA, enabling reusable code, independent secrets, and unit testing across Actions.

└──▷ **HOW TO FIND IT**

Create a new Action Module in the Auth0 Dashboard to house shared initialization logic (e.g. an SDK client) that multiple Actions can import instead of each re-running setup code.

> 📍 In the Auth0 Dashboard, go to Actions › Modules › Create Module, give it a name, add any required secrets and npm dependencies, then write and publish your shared code.

Adds `@auth0/actions` npm package providing unit-testing helpers and IntelliSense (autocompletion and type hints) for Actions Modules in any editor.

Introduces Actions Modules — create, manage, and share reusable code across multiple Actions and Triggers within an Auth0 Tenant.

Supports independent secrets and dependencies defined per Action Module, decoupled from the Actions that import them.

Enables cross-trigger reuse of custom functions by importing a single Module into any Action regardless of which Trigger it is bound to.

Improves execution performance by moving expensive initialization logic into a Module that is initialized once and reused across Action invocations.

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ

Clean up unapproved password manager use with a single batched nudge NOTES ↗

RANK

Discover every SaaS and AI asset, harden security posture, and scale identity governance.

> Nudge Security added a bulk nudge that consolidates password manager compliance messaging into one batched notification per employee.

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

01 **Bulk nudge for approved password manager use** NEW 85

Employees using unapproved password managers can now be nudged in bulk: filter the accounts inventory by password manager observed at login, multi-select the flagged accounts, and send the Use approved password manager nudge. Each employee receives a single batched nudge listing all their affected accounts instead of one notification per account.

Identify and notify all employees using unapproved password managers in one pass, without generating per-account notification noise.

> 📍 In the Nudge Security console, go to Accounts Inventory › Filter by Password Manager (select all not marked as approved) › multi-select the resulting accounts › Send nudge › Use approved password manager .

— Names exact UI steps and nudge type; clear before/after on messaging volume.

Clean up unapproved password manager us…

Clean up unapproved password manager use with a single batched nudge NOTES ↗

LATEST STABLE

Nudge Security adds bulk Use approved password manager nudge — one batched message per employee across all flagged accounts.

└──▷ **HOW TO FIND IT**

Identify and notify all employees using unapproved password managers in one pass, without generating per-account notification noise.

> 📍 In the Nudge Security console, go to Accounts Inventory › Filter by Password Manager (select all not marked as approved) › multi-select the resulting accounts › Send nudge › Use approved password manager .

WAS THIS USEFUL?

Clerk

Feedback ⓘ Trace SMS delivery in Application Logs (2026-09-17)

BLOG ↗ RANK

Clerk is an authentication and user-management platform that provides pre-built UI components, APIs, and SDKs for adding sign-up, sign-in, and identity features to applications.

> Clerk added SMS delivery lifecycle tracking to Application Logs, letting teams trace the status of every Clerk-delivered SMS.

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

01 **SMS delivery tracing in Application Logs** NEW 40

Clerk Application Logs now expose the full delivery lifecycle of every SMS message sent through Clerk, letting teams trace the status of each Clerk-delivered SMS.

— Names the log feature but no fields, statuses, or navigation path

Trace SMS delivery in Application Logs …

> Trace SMS delivery in Application Logs (2026-09-17) BLOG ↗ LATEST STABLE
>
> Clerk Application Logs now expose the full delivery lifecycle of every SMS message sent through Clerk.
>
> Adds SMS delivery lifecycle tracking to Application Logs, letting teams trace the status of every Clerk-delivered SMS.

WAS THIS USEFUL?

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ

1.28.574 - Richer session timelines and inline video artifacts (2026-09-18)

BLOG ↗ RANK

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

> Runlayer added Microsoft Teams bot provisioning across both the self-hosted Operator and the Terraform module, introduced Agent Accounts giving each agent its own identity with independent policies and audit trail, and rounded out the session/run view with skipped-message timelines, inline video playback, and tighter audit-log redaction.

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

01 Microsoft Teams integration in Terraform module

BREAKING

95

The Terraform module adds Microsoft Teams integration configured via `MS_TEAMS_APP_CLIENT_ID`, `MS_TEAMS_RUNLAYER_TENANT_ID`, `MS_TEAMS_AZURE_SUBSCRIPTION_ID`, `MS_TEAMS_AZURE_RESOURCE_GROUP`, and `MS_TEAMS_APP_URL`, with the client secret supplied via `ms_teams_app_client_secret` (recommended via `TF_VAR_ms_teams_app_client_secret`); the integration is off unless all five values are set. Provisioning creates a per-agent Entra app registration and Azure Bot Service resource, requiring a multi-tenant app registration with Graph permissions `Application.ReadWrite.All`, `AppCatalog.ReadWrite.All`, and `TeamsAppInstallation.ReadWriteForUser.All` (admin-consented), a resource group granting Azure Bot Service Contributor, and the `Microsoft.BotService` provider registered on the subscription. The Redirect URI `https://<domain_name>/api/v1/integrations/ms-teams/admin-consent-callback` must be registered under Authentication → Web on the app registration or admin consent fails with `AADSTS500113`. Breaking: the domain name set at agent provisioning time is effectively immutable — each agent's Azure Bot resource stores the messaging endpoint captured at creation and never rewrites it, so changing the domain orphans existing bots until each is deleted and recreated.

— Exact Graph permissions, redirect URI, and breaking domain caveat named product docs

02 Microsoft Teams provisioning via Runlayer Operator NEW

90

The Runlayer Operator gains six environment variables to enable Microsoft Teams bot provisioning: `MS_TEAMS_APP_CLIENT_ID`, `MS_TEAMS_APP_CLIENT_SECRET`, `MS_TEAMS_RUNLAYER_TENANT_ID`, `MS_TEAMS_AZURE_SUBSCRIPTION_ID`, `MS_TEAMS_AZURE_RESOURCE_GROUP`, and `MS_TEAMS_APP_URL`. `MS_TEAMS_APP_URL` sets the public HTTPS base for the Bot Framework messaging endpoint and admin-consent callback, defaults to the operator's own host, and must match the Redirect URI registered on the Entra app or admin consent fails with `AADSTS500113`; Teams provisioning stays off unless all six variables are set.

Enable Microsoft Teams bot provisioning on a self-hosted Runlayer Operator deployment by supplying all required `MS_TEAMS_*` variables — omitting any of them leaves Teams provisioning off.

```text
MS_TEAMS_APP_CLIENT_ID=<entra-app-client-id>
MS_TEAMS_APP_CLIENT_SECRET=<entra-app-client-secret>
MS_TEAMS_RUNLAYER_TENANT_ID=<entra-tenant-id>
MS_TEAMS_AZURE_SUBSCRIPTION_ID=<azure-subscription-id>
MS_TEAMS_AZURE_RESOURCE_GROUP=<resource-group-name>
MS_TEAMS_APP_URL=https://runlayer.example.com
```

— Full env var list, default behavior, and exact error code given `product docs`

03 **Session and run view improvements** IMPROVED 60

The Agents Sessions list now shows Slack messages the agent chose to skip alongside runs, with search covering both skipped messages and runs; video artifacts play inline in the run view with the artifact size limit raised to 25 MB; and clicking a trace header with a single tool call now opens its details directly.

— UI locations and 25 MB limit named, no config keys

`1.28.574 - Richer session timelines and…`

THINNER COVERAGE BELOW

04 **Agent Accounts for per-agent identity** NEW 45

Adds Agent Accounts so each agent gets its own identity in Runlayer, with independent policies, delegations, tokens, and audit trail.

— Describes capability but no config or API surface named `product docs`

05 **Managed device login host badge** NEW 45

Managed devices show a badge when a user's login host differs from the MDM-managed host, and hooks target the managed host.

— Describes trigger condition but no exact UI path

`1.28.574 - Richer session timelines and…`

06 **OAuth token redaction in audit logs** IMPROVED 25

OAuth token keys are redacted in audit log details.

— One-line description with no mechanism or scope given

1.28.574 - Richer session timelines and…

└──▷ **BREAKING ON UPGRADE**

! The domain name configured at agent provisioning time is effectively immutable: each agent's Azure Bot resource stores the messaging endpoint captured at creation time and is never rewritten, so changing the domain orphans existing bots until each is deleted and recreated.

1.28.574 - Richer session timelines and inline video artifacts (2026-09-18)

BLOG ↗ LATEST STABLE

Runlayer 1.28.574 adds skipped-message timelines, inline video artifacts (25 MB), login-host badges, and OAuth token redaction in audit logs.

Agents Sessions list now shows Slack messages the agent chose to skip alongside runs, with search covering both skipped messages and runs.

Video artifacts play inline in the run view, and the artifact size limit is raised to 25 MB.

Clicking a trace header with a single tool call opens its details directly.

Managed devices show a badge when a user's login host differs from the MDM-managed host, and hooks target the managed host.

OAuth token keys are redacted in audit log details.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ToolHive

Feedback ⓘ v0.50.0 NOTES ↗ RANK

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

> ToolHive v0.50.0 strengthens the auth server with RFC 8707 resource indicators, RFC 8628 device authorization flow, and ID-JAG assertion support, while also hardening plugin installs with catalog provenance checks, moving secrets encryption to Argon2id key derivation, and adding key rotation workflows.

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

01 **RFC 8707 resource indicators for token requests** NEW 55

Adds RFC 8707 resource indicator parameters to the auth server's token-request handling, enabling audience-bound token requests.

– Names RFC and mechanism but no exact endpoint or flag v0.50.0

02 **RFC 8628 device authorization grant** NEW 55

Adds RFC 8628 device authorization grant (device flow) support, including a device flow verification page and Dynamic Client Registration (DCR) support.

– Describes flow and page but no navigation or command given v0.50.0

03 **ID-JAG assertion support in auth server** NEW 55

Adds acceptance of ID-JAG assertions via a bound JWT-bearer handler in the auth server.

– Specific mechanism named but no config or command shown v0.50.0

04 **Argon2id secrets key derivation** IMPROVED 50

Derives the secrets encryption key using Argon2id, replacing the previous key derivation approach.

– Names algorithm and change but no config key or migration steps v0.50.0

05 **Catalog provenance enforcement for plugins** NEW 45

Enforces catalog provenance for plugin installs, ensuring plugins are verified against registry provenance before installation.

– Explains enforcement but no install command or flag named v0.50.0

06 Lock-driven key rotation for skills and plugins NEW

30

Adds lock-driven skill key rotation and plugin key rotation workflows.

— Thin description with no mechanism or command detail v0.50.0

v0.50.0 NOTES ↗ LATEST STABLE

ToolHive v0.50.0 adds RFC 8707 resource indicators, RFC 8628 device authorization grant, ID-JAG assertion support, catalog provenance enforcement, and Argon2id secrets key derivation.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.50.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.50.0
```

Adds RFC 8707 resource indicator parameters to the auth server's token-request handling, enabling audience-bound token requests.

Adds RFC 8628 device authorization grant (device flow) support, including a device flow verification page and Dynamic Client Registration (DCR) support.

Adds acceptance of ID-JAG assertions via a bound JWT-bearer handler in the auth server.

Enforces catalog provenance for plugin installs, ensuring plugins are verified against registry provenance before installation.

Derives the secrets encryption key using Argon2id, replacing the previous key derivation approach.

+1 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 19, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue