THE CYBER TOOLCHAIN

NO. 123

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 28, 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 242 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 74 tools

▾ // THE BRIEF HIDE

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

Three unrelated things stand out. StepSecurity now inventories Claude Code and Codex plugins across machines and repos, so you finally which agents your developers installed. Socket tells you whether

vim ? help

a malware campaign actually touched your org, not just that it exists. Censys made its hunting workflow an API, so SOAR playbooks can submit indicators and pull back campaign clusters.

> IDENTIFY
>
> **Find out which AI coding agents your developers actually installed**
>
> StepSecurity enumerates Claude Code and Codex plugins org-wide across developer machines and GitHub repos; Castle exposes a structured `agent` object in its API response so teams can see and gate agent traffic at the app edge. Both close the same blind spot from opposite ends: agents arrive without a procurement step, and until now nobody had a list.
>
> StepSecurity · Castle

> RESPOND
>
> **Skip the question of whether a published campaign actually touched you**
>
> Socket's Threat Intel Campaigns page now resolves per-package impact down to the originating scan or Firewall event, turning a public advisory into a yes/no about your own estate instead of a manual dependency sweep. Censys exposed its Investigator hunt as an API, so a playbook can submit indicators, push raw intel documents, and retrieve campaign-cluster reports without an analyst in the console.
>
> Socket · Censys

> PROTECT
>
> **Let the scanner propose and validate the fix, not just file the finding**
>
> SonarQube's Remediation Agent now covers Hunter Agent findings — broken access control, business logic flaws, auth gaps — with CI validation of the generated fix, which is the step that normally keeps auto-fix out of production. GitGuardian cut scan noise from IDE history and build/cache folders and expanded validity checks to 501 secret types, so triage starts with secrets that still work.
>
> SonarQube · GitGuardian SaaS

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

ProjectDiscovery naabu

Feedback ⓘ v2.7.0 NOTES ↗

RANK

A fast port scanner written in go with a focus on reliability and simplicity. Designed to be used in combination with other tools for attack surface discovery in bug bounties and pentests

> naabu v2.7.0 adds native service version fingerprinting, protocol-aware UDP probes, dead-host reporting during discovery, plus internal performance gains and musl release builds.

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

01 Native service version fingerprinting with -sV NEW

88

Adds `-sV` / `-service-version` flag for service version detection using an embedded nmap-service-probes database, with `-sV-fast` to limit probing to port-hinted services only, `-sV-timeout` to control probe timeout, `-sV-workers` to set concurrency, and `-sV-probes` to supply a custom nmap-service-probes file path.

Fingerprint services on the top 1000 ports of a target, using the embedded nmap probes and limiting workers for a controlled scan.

```
$ naabu -host example.com -top-ports 1000 -sV -sV-workers 10 -sV-timeout 10s -json -o results.json
```

— Names all flags with mechanism and a runnable example v2.7.0

02 Protocol-specific UDP probes with -uP NEW

72

Adds `-uP` / `-udp-probes` flag to send protocol-specific payloads on UDP scans, using the nmap-service-probes database to identify UDP services.

Run a UDP scan with protocol-specific payloads to identify UDP services on common ports.

```
$ naabu -host 192.168.1.0/24 -p 53,161,500,4500 -s SYN -uP -json
```

— Named flag and mechanism with runnable example v2.7.0

03 Dead host reporting with -show-dead NEW

65

Adds `-show-dead` flag to report hosts that did not respond during host discovery; requires host discovery to be enabled to surface asset inventory gaps.

Discover live hosts in a CIDR and surface dead hosts that did not respond, useful for asset inventory gaps.

```
$ naabu -host 10.0.0.0/24 -sn -show-dead -silent
```

— Named flag with precondition and runnable example v2.7.0

THINNER COVERAGE BELOW

04 **Musl release archives for Alpine Linux** NEW 30

Adds musl release archives for Alpine Linux.

— Names platform but no further detail v2.7.0

05 **Internal scanning performance optimizations** IMPROVED 11

Delivers faster scanning performance via internal optimizations.

— No mechanism or numbers given v2.7.0

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

xalgorix

Feedback ⓘ v4.6.173 NOTES ↗ RANK

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

> xalgorix fixed a parser gap where DSML-namespaced tool-call markup from models was going unrecognized, causing scans to stall.

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

01 **DSML-namespaced tool-call markup parsing** IMPROVED 57

Added parser support for DSML-namespaced tool-call markup (`< | DSML | invoke name="…">` / `< | DSML | parameter …>`) so model invocations in this format are no longer silently dropped, preventing scans from stalling at the `no tool call` stop limit when a model emits DSML-tagged invocations.

— Names exact markup syntax and bug fixed, but no usage example v4.6.173

NEW INTEGRATION GPT-6.1 Terra model support

└──▷ ALSO FROM THESE RELEASES

Add both Refuseless and Abliteration providers to your gateway for red-teaming or security research workflows that require uncensored model responses.

yaml

```yaml
model_list:
  - model_name: refuseless-chat
    litellm_params:
      model: refuseless/glm-5.3-abliterated
      api_key: os.environ/REFUSELESS_API_KEY
  - model_name: abliteration-chat
    litellm_params:
      model: abliteration/abliterated-model-large-v2
      api_key: os.environ/ABLITERATION_API_KEY
```

Send a chat completion to the Refuseless abliterated GLM model via the LiteLLM gateway, using an existing virtual key.

```
$ curl "$LITELLM_BASE_URL/v1/chat/completions" \
  -H "Authorization: Bearer $LITELLM_API_KEY" \
  -H "Content-Type: application/json" \
  -d '{
    "model": "refuseless-chat",
    "messages": [{"role": "user", "content": "Explain how DNS
works."}],
    "reasoning_effort": "low"
  }'
```

announced · Oct 26, 2026 BLOG ↗

LiteLLM gains Day 0 GPT-6.1 Terra support with automatic Responses API routing for tool calls, a new `POST /reload/model_cost_map` endpoint, and `reasoning_effort` parameter support.

Adds `openai/gpt-6.1-terra` as a supported model in `config.yaml` under `model_list`, available on v1.76.0 and above.

Automatically routes `/chat/completions` requests with tools to `/v1/responses` when reasoning is enabled for GPT-6.1 Terra.

Adds `POST /reload/model_cost_map` endpoint (or UI Reload Price Data under Models + Endpoints) to pull updated pricing without a proxy restart.

Supports the `reasoning_effort` parameter (values: `none`, `low`, `medium`, `max`; default `medium`) for GPT-6.1 Terra via both the proxy and the Python SDK.

WAS THIS USEFUL? 👍 👎

Dreadnode

Feedback ⓘ New navigation now live for all users

NOTES ↗

Documentation for the Dreadnode AI security platform — self-hosting guides, platform concepts, SDK reference, and CLI documentation.

Dreadnode adds a `tool_call_attack_landed` scorer for judge-free agentic red teaming and ships a third-party notices file for air-gapped Enterprise installs.

NEW DETECTION Adds `tool_call_attack_landed` scorer for agentic red teaming — measures whether an agent emitted a specific attacker-intended tool call (name and arguments), providing a deterministic, judge-free attack success rate signal.

Enterprise installs now include a third-party notices file downloadable from Account Settings or at `https://<domain>/third-party-notices`, air-gap compatible and version-matched.

Redesigned navigation is now the default for all users on SaaS and Enterprise.

NEW DETECTION New structural and long-context jailbreak attacks in the airt SDK

NEW DETECTION New classifier-evasion attacks in the airt SDK

WAS THIS USEFUL? 👍 👎

◆ Password & Credential Cracking

seclists

Feedback ⓘ 2026.2 NOTES ↗

SecLists is the security tester's companion. It's a collection of multiple types of lists used during security assessments, collected in one place.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2026.2
https://github.com/danielmiessler/SecLists.git
# already have the repo? check out this version:
$ git checkout 2026.2
```

NEW DETECTION Claude files added to quickhits.txt

NEW DETECTION LureScope observed SSH credentials wordlist

NEW DETECTION Risk-classified SQLi payload list

NEW DETECTION Template engine fuzzing wordlist

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Software Supply Chain Security

Cloudsmith

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗

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

CLI v1.30.0 adds TLS-intercepting proxy support and anonymous error reporting; npm provenance verification and install performance improved; exact usage-limit inputs added.

Adds TLS-intercepting proxy support to the Cloudsmith CLI v1.30.0, enabling connections through TLS-intercepting proxies without disabling certificate verification.

Cloudsmith's npm registry now serves npm attestations and signing-key APIs, enabling npm provenance verification through Cloudsmith just as it works against the upstream registry.

Workspace owners and managers can now type exact numerical values for usage limits on the workspace Settings tab under `Usage limits`, replacing imprecise sliders.

CLI v1.30.0 now sends anonymous error reports to Cloudsmith by default so crashes can be identified without a support ticket.

Reduces pnpm client request count against Cloudsmith and reduces upstream lookups per npm repository, speeding up npm installs with no configuration changes required.

NEW SUPPORTED PLATFORM Mumbai storage region

└──▷ BREAKING ON UPGRADE

! Username and password login via `cloudsmith login` and `cloudsmith token` stops working on November 20, 2026 — scripts and CI jobs using these commands must switch to API key authentication before that date.

WAS THIS USEFUL? 👍 👎

plumber

Feedback ⓘ v0.7.16 NOTES ↗ RANK

Plumber detects CI/CD security issues in your GitHub workflows and gives you a score

> Plumber's `platform push` command now uploads the GitHub bill of materials alongside its existing payload.

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

01 **GitHub bill of materials in platform push** IMPROVED

50

The `platform push` command now includes the GitHub bill of materials in its upload payload.

— Names exact command and payload addition, lacks further mechanism detail. v0.7.16

WAS THIS USEFUL? 👍 👎

Endor Labs

Feedback ⓘ

Introducing AURI Fix for Open Source | Endor Labs (2026-10-27) BLOG ↗ RANK

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

> Endor Labs launched AURI Fix for Open Source (private beta), a new `endorctl fix` CLI-triggered remediation system that can patch vulnerabilities via code mitigations, configuration changes, or fully agentic PR workflows.

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

01 AURI Fix for Open Source via `endorctl fix` NEW 80

Endor Labs introduced AURI Fix for Open Source, a private-beta remediation system triggered by the new `endorctl fix` CLI command, which automatically assesses a reachable finding and opens a pull request with the safest fix — an upgrade, a first-party code mitigation, a targeted dependency patch, a configuration change, or dependency replacement.

Trigger AURI Fix from the CLI to automatically assess a reachable finding and open a pull request with the safest remediation — upgrade, code mitigation, or config change.

```
$ endorctl fix
```

— Names the exact CLI command and lists all remediation options launch-20261027-24c318ac Introducing AURI Fix for Open Source | …

THINNER COVERAGE BELOW

02 Configuration-based mitigations NEW 53

AURI identifies and sets the specific property or configuration flag that disables a vulnerable code path, re-evaluates the configuration to confirm no exploitable path remains, and opens a pull request with the change.

— Describes full mechanism from detection to PR but no named flag

launch-20261027-24c318ac Introducing AURI Fix for Open Source | …

03 Agentic and policy-based remediation workflow NEW 53

A policy-based workflow lets the Endor Labs policy engine automatically open a pull request for a fix, while an agentic workflow has an AURI agent triage findings, select the remediation method (upgrade, code mitigation, config change, patch, or replacement), apply the fix, and open the pull request for human approval.

— Describes two distinct automation paths without a named interface

launch-20261027-24c318ac | Introducing AURI Fix for Open Source | …

04 **Application-side code mitigations** NEW 50

AURI can change how an application calls a vulnerable dependency to break the exploit path, typically in fewer than 10 lines of code, without waiting for an upstream patch to be released.

— Explains mechanism and gives a concrete size bound

Introducing AURI Fix for Open Source | …

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ announced · Oct 27, 2026 BLOG ↗ RANK

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

> StepSecurity shipped org-wide inventory of AI agent plugins (Claude Code, Codex) across developer machines and GitHub repos, overhauled Dev Machine Guard's malicious-file detection with dedicated Rules and Findings pages, and launched a beta Malicious Files scanner for GitHub repositories covering known supply-chain campaigns like Shai-Hulud, GlassWorm, and Miasma.

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

01 Malicious Files Rules and Findings pages for Dev Machine Guard

NEW 91

Renames Suspicious Files to Malicious Files and adds a Rules page under Developer Machines listing every detection rule with its campaign, indicator groups, and file patterns, searchable by rule name and filterable by campaign, with per-rule activate/deactivate controls that stop scanning across all enrolled devices and move existing findings to a reversible Suppressed state. A new Findings page under Developer Machines groups flagged files into New, Resolved, and Suppressed tabs with device, user, OS, campaign, confidence, and status metadata, filterable by confidence or campaign and exportable to CSV.

Suppress noisy detections fleet-wide by deactivating a rule from the Rules page, keeping findings visible under the Suppressed tab without deleting them.

> 📍 In Dev Machine Guard, go to Developer Machines › Rules, locate the rule by name or filter by campaign, then click Deactivate to stop it scanning every enrolled device and move its existing findings to the Suppressed tab. Click Activate to re-enable it when needed.

Triage high-confidence malicious file hits across your fleet and export them for further analysis or ticketing.

> 📍 In Dev Machine Guard, go to Developer Machines › Findings, set the confidence filter to high, select the New tab to see unflagged items, then click Export to download a CSV of the filtered findings.

— Exact UI navigation and toggle behaviour, no API surface named.

Dev Machine Guard: Malicious Files Rule…

02 AI agent plugin inventory for Claude Code and Codex NEW

90

Dev Machine Guard gains a plugin detail panel listing marketplace source, plugin file origin, per-device installation scope (User, Project, Local, System), version, state, and component breakdown by kind (skills, commands, MCP servers), with Claude Code skill usage counts; plugin-sourced skills and MCP servers are attributed back to their originating plugin on the existing Agent Skills and MCP Servers pages. A new Agent Plugins page under GitHub > AI Agents inventories plugin manifests (`.claude-plugin/plugin.json` for Claude Code, `.codex-plugin/plugin.json` for Codex) published across org repositories, showing publishing repository, plugin folder path, declared metadata, branch, and commit scanned, plus a Used by view listing every repository whose settings enable the plugin for all contributors. Supports CSV export of the Dev Machine Guard inventory, and filtering by installation scope/source kind/enablement (Dev Machine Guard) or by agent/repository (GitHub).

Identify which developer machines have a specific Claude Code plugin enabled and where its files actually came from — useful when investigating a suspicious marketplace listing.

> 📍 In Dev Machine Guard, go to IDE & AI Agents › Agent Plugins, search by plugin name or publisher, filter by source kind (e.g., `github` or `local`), and select the plugin to open its detail panel showing per-device installation scope, version, state, and marketplace vs. plugin file source.

Determine which repositories in your org enable a published Claude Code plugin for all contributors — catching broad exposure before a compromised plugin update rolls out.

> 📍 In StepSecurity, go to GitHub › AI Agents › Agent Plugins, filter by agent (Claude Code), select the plugin, and review the Used by panel listing every repository whose `.claude/settings.json` enables the plugin.

Export the full org-wide AI agent plugin inventory as CSV for offline review or SIEM ingestion.

> 📍 In Dev Machine Guard, go to IDE & AI Agents › Agent Plugins and click the CSV export button to download the full inventory including plugin names, publishers, source kinds, versions, and enabled/disabled device counts.

— Names exact manifest paths and UI surfaces, but no API endpoint given.

launch-20261027-75c6326a

03 Malicious Files scanning for GitHub repositories (Beta) NEW

90

Adds a Malicious Files section under GitHub scans that inspects full repository history — every branch, pull request, and tag — for files planted by known supply chain attack campaigns such as Shai-Hulud, GlassWorm, and Miasma. Includes a Rules view showing each rule's campaign, severity, revision, and ordered checks; a Scan Status view showing the most recent scan per repository with commits/file versions scanned and last/next scan times; and a Findings view listing matched files with branches, pull requests, tags, matched checks, and the introducing commit, filterable by severity, campaign, or rule, with CSV export.

Review active detection rules to understand which supply chain campaigns are covered and at what severity before interpreting findings.

> 📍 In the StepSecurity console, go to GitHub › Malicious Files › Rules to see each rule's campaign, severity, revision, and the ordered checks applied to every file version.

Export all Findings for a repository to CSV for offline triage or ingestion into a SIEM, filtered to high-severity results.

> 📍 In the StepSecurity console, go to GitHub › Malicious Files › Findings, filter by severity, then select Export to CSV.

— Names specific campaigns and views; UI-only, no API endpoint given.

Malicious Files for GitHub Repositories…

WAS THIS USEFUL? 👍 👎

sbom-operator

Feedback ⓘ 0.48.0 NOTES ↗ RANK

Catalogue all images of a Kubernetes cluster to multiple targets with Syft

> sbom-operator 0.48.0 adds global per-pod parent project configuration for its Dependency-Track integration.

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

01 **Global per-pod parent project config for Dependency-Track**

IMPROVED 40

Adds global configuration per pod for parent projects in the Dependency-Track (`dtrack`) integration.

— Names the config area but no exact key or flag given 0.48.0

WAS THIS USEFUL? 👍 👎

Socket

Feedback ⓘ announced · Oct 26, 2026 BLOG ↗ RANK

Socket is a supply chain security platform that analyzes open-source package dependencies (npm, PyPI, etc.) to detect malicious code, risky behaviors, and vulnerabilities before developers install them.

> Socket's Threat Intel Campaigns page now shows whether a malware campaign actually hit your organization, with drill-down from per-package impact straight to the originating scan or Firewall event.

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

01 **Campaign impact tracking in Threat Intel** NEW 75

The Threat Intel → Campaigns page now shows a per-organization status badge on every malware campaign — Blocked by Socket, Detected in Firewall/SCA, or Not Detected — computed from repository scans and Firewall event traffic. A new Org Impact tab on each campaign lists each affected package and version separately along with the repositories, branches, and first/last-seen dates, which link directly to the originating scan report or matching Firewall event; new Blog and Affected Packages tabs add Socket write-ups and the full list of affected packages to the campaign detail view.

Determine whether a newly publicized malware campaign reached your org and whether Socket blocked or only detected it, then trace every affected repository and branch to the exact scan.

> 📍 In the Socket dashboard, go to Threat Intel → Campaigns. Locate the campaign in question and check the org-level status badge (Blocked by Socket, Detected in Firewall/SCA, or Not Detected). Click the campaign and open the Org Impact tab to see each affected package version, the repositories and branches that contained it, and first/last-seen dates. Click any date to open the corresponding scan report or Firewall event.

— Names exact statuses, tabs and page path with clear navigation steps.

launch-20261026-420a0dd6

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Openlayer

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗

Openlayer is the AI governance and observability platform for regulated enterprises: one place to test, monitor, and govern every model, agent, and AI system your company runs.

Openlayer adds governance lifecycle stages, per-step test evaluation, EU data residency, Gateway routing/allow-block lists, new API endpoints, and multiple project dashboards.

Adds `evaluationScope: step` config field for per-step evaluation of tests on each trace step.

Adds `openlayer tests delete` CLI command — archives by default, `--force` permanently deletes.

Extends `openlayer init` with Copilot CLI and Gemini CLI support, custom CA bundles, a 'debug with an agent' option, and `.openlayerignore` for smaller uploads.

Adds Gateway `GET /v1/models` (per key) and OpenAI-compatible `POST /v1/embeddings` API endpoints.

Adds governance registry and framework endpoints to the public API and SDKs, writable with a workspace API key.

+26 more

NEW COMPLIANCE FRAMEWORK NYDFS Circular Letter No. 7 (2024) built-in framework

NEW DETECTION Bias and fairness tests for protected group disparities

NEW DETECTION Prompt injection tests with multilingual support

NEW DETECTION CPNI and telecom guardrail with PII recognizers

NEW INTEGRATION Databricks AI/BI Genie discovery and monitoring integration

NEW INTEGRATION Palantir Foundry AIP Agents and AIP Logic integration

NEW INTEGRATION Claude Code and Claude Desktop OpenTelemetry telemetry

NEW INTEGRATION AI usage-log data sources: OpenAI Compliance Logs, Claude Compliance API, Microsoft 365 Copilot

NEW INTEGRATION CrowdStrike Falcon Guardian integration for AI coding agents

NEW INTEGRATION Microsoft Agent Framework tracing integration

UPDATED INTEGRATION Google ADK tracing updated for agent handoffs and google-adk 2.x

NEW INTEGRATION Remote MCP connector with OAuth2 and permission levels

NEW SUPPORTED PLATFORM AWS ECS Fargate reference deployments with Terraform

└──▷ **BREAKING ON UPGRADE**

! Gateway `Teams` are renamed to `Groups` — any configuration, API calls, or documentation referencing `Teams` in the Gateway admin UI must be updated to `Groups`.

WAS THIS USEFUL?  

OpenQodex

Feedback ⓘ openqodex@0.16.0 NOTES ↗ RANK

Open source code review that runs in your coding agent before you push. Scanners plus a separate AI reviewer, on your own login, no key.

> OpenQodex 0.16.0 deepens its Java code graph with test-to-code linking, richer class and pattern-variable modeling, and extends chained-call tracking to Go and TypeScript, alongside a tweak to finding report formatting.

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

01 **Java code graph models anonymous, local and enum classes**

IMPROVED 63

The Java code graph now reads anonymous classes, local classes, and enum constant bodies as first-class definitions, with scoped local-class names and qualified `this` binding.

— Names concrete Java constructs and binding behavior openqodex@0.16.0

THINNER COVERAGE BELOW

02 **Interface default-method resolution improved** IMPROVED 58

Java interface default-method resolution now selects zero-parameter defaults from the most specific interface, while parameterized methods keep both interface declarations as possible candidates.

— Explains before/after resolution logic clearly openqodex@0.16.0

03 **Finding `problem` field allows longer prose** IMPROVED 50

A finding's `problem` field can now contain up to three sentences, and the same prose checks applied to all other fields are now applied to finder claims as well.

— Names the field and limit but minimal mechanism openqodex@0.16.0

04 **Pattern variable tracking through negated guards** IMPROVED 45

OpenQodex follows Java pattern variables through negated guards that exit and through negated while loops.

— Specific mechanism but no concrete usage path openqodex@0.16.0

05 **JUnit tests linked to Java call graph** NEW 41

OpenQodex now links JUnit 4 and JUnit 5 test methods to the Java code they call within the code graph, connecting test coverage to the underlying call chain.

— Describes what links but no mechanism or scope detail `openqodex@0.16.0`

06 **Chained-call tracking for Go and TypeScript** IMPROVED

41

Code graph chained-call tracking is extended to typed Go and TypeScript receivers.

— Names languages but little mechanism detail `openqodex@0.16.0`

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

GitHub Copilot supports server-to-server authentication (2026-10-27) BLOG ↗

RANK

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

> GitHub Copilot added public-preview server-to-server authentication for GitHub Apps, letting automated agents call the Copilot SDK under an organization identity without a signed-in user.

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

01 **Server-to-server authentication for Copilot SDK** NEW

81

GitHub Apps installed on organizations can now authenticate to the Copilot SDK using a short-lived installation access token instead of a user's credentials (public preview), enabling agents and automated services to call Copilot under an org identity. A new read-only Copilot Requests organization permission scopes what an app can do when authenticating as this workload identity, and app-driven usage is attributed to the installing organization — including model policies, budgets, and spending controls — with app traffic identifiable in AI usage CSV reports via the `app_client_id` field.

— Names permission, token mechanism and CSV field but no setup steps.

GitHub Copilot supports server-to-serve…

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Integration updates archive (2026-10-27)

BLOG ↗ RANK

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

> Veracode's largest update this window is a refreshed CLI (v2.53.0) with new `veracode scan static`, `veracode findings static`, `veracode sca`, and `auth login` commands plus a wide set of `veracode package` flags and environment variables, alongside a new COBOL scanner, `--include` module filtering in Pipeline Scan, a GitHub Workflow App for org-wide scanning, a JetBrains SCA plugin, and the end of TFS support in the Azure DevOps Extension.

└──▷ **WHAT SHIPPED · 20 FEATURES** most completely described first ? what's the number?

01 **New CLI commands for static scans and SBOM** NEW 90

Veracode CLI 2.53.0 adds `veracode scan static` for running static scans and reviewing security findings within a Git workflow, `veracode findings static` for accessing cached static scan findings in the local Veracode directory, `veracode sca` for SBOM generation via SCA agent-based scans publishing results in CycloneDX 1.6 and SPDX 2.3 formats, and `auth login` for onboarding new users via OAuth2.

Run a static scan inside a Git repo and immediately review findings without leaving the CLI — ideal for gating a PR before merge.

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

— Names every command and format, plus a runnable example. CLI updates (2026-10-27)

02 **Packaging flags and environment variables in `veracode package`**

NEW 85

`veracode package` adds the `--csharp-source` flag to package C# source code directly for Static Analysis without compiling it into binaries, the `--skip-packagers` flag to exclude specific artifacts (such as unsupported SQL) from packaging without affecting other languages, the `--exclude` flag to prevent specified Java packages or JAR files from being scanned, the `VERACODE_PACKAGE_GOLANG_GENERATE` environment variable to enable client-side CGO code generation when packaging Go applications, the `SRCLR_MAKE_WINDOWS` environment variable to preprocess C/C++ source files with Make on Windows using a native Windows gcc or clang compiler, the `--java-source` flag and `VERACODE_PACKAGE_JAVA_SOURCE=true` environment variable to enable Java source code packaging, support for `requirements-dev.txt` in

Python packaging workflows, support for the Poetry package management tool for Python autopackaging, automatic inclusion of `.yaml`, `.yml`, and `.properties` files when packaging Java source code, support for packaging Go v1.24, and support for a generic packager for Apex.

— Enumerates every flag and env var verbatim, no single run example.

CLI updates (2026-10-27)

03 **.NET cleanser deprecations for several CWEs** DEPRECATED

80

Veracode deprecates .NET cleansers for flaw classes CWE-93, CWE-113, and CWE-117 across `antixsslibrary.dll: Microsoft.Security.Application.AntiXss.HtmlAttributeEncode`, `antixsslibrary.dll: Microsoft.Security.Application.AntiXssEncoder.HtmlAttributeEncode`, `antixsslibrary.dll: Microsoft.Security.Application.Encoder.HtmlAttributeEncode`, `antixsslibrary.dll: Microsoft.Security.Application.Encoder.HtmlEncode`, `mscorlib.dll: System.Security.SecurityElement.Escape`, `system.dll: System.Net.WebUtility.HtmlEncode`, `system.web.dll: System.Web.HttpServerUtility.HtmlEncode`, `system.web.dll: System.Web.Security.AntiXss.AntiXssEncoder.HtmlEncode`, `system.web.dll: System.Web.Util.HttpEncoder.HtmlAttributeEncode`, `system.web.dll: System.Web.Util.HttpEncoder.HtmlEncode`, `system.web.mvc.dll: System.Web.Mvc.HtmlHelper.AttributeEncode`, `system.web.mvc.dll: System.Web.Mvc.HtmlHelper.Encode`, `system.windows.browser.dll: System.Windows.Browser.HttpUtility.HtmlEncode`, `system.windows.dll: System.Net.HttpUtility.HtmlEncode`, and `System.Runtime.dll: System.Net.WebUtility.HtmlEncode`; applications relying on these as cleansers will no longer have those flaws suppressed.

— Names every deprecated function and the resulting impact.

SAST updates archive (2026-10-27)

04 **API Wrapper enhancements** NEW

80

The Java API Wrapper adds support for the `-includenewmodules` parameter and enables sending requests to the Identity REST API from within the wrapper, while Veracode C# API Wrapper 23.5.8.8 adds proxy support via the `https_proxy` environment variable, accepting host, port, and credentials.

Route all Veracode C# API Wrapper calls through a corporate proxy without modifying wrapper arguments.

```
$ export
https_proxy=http://proxyuser:proxypass@proxy.corp.example.com:808
0
```

— Names the flag and API but no usage example given. Includes exact env var and a runnable export command.

Integration updates archive (2026-10-27)

05 **Toolchain discovery in `veracode package discover`** IMPROVED

75

`veracode package discover` detects toolchains required to create scannable artifacts, including JavaScript projects using npm or Yarn and Java projects using the Gradle package manager, and introduces exit codes 5, 6, and 7 to improve CI/CD pipeline automation and error handling.

— Names package managers and exit codes but no sample invocation.

CLI updates (2026-10-27)

06 **`veracode scan` findings and policy enhancements** IMPROVED

75

`veracode scan` now categorizes secret-type findings as Critical, prioritizes vulnerability results from the NVD database when available, enables applying a custom policy defined in the Veracode Platform and viewing policy evaluation for individual findings, enables specifying a custom location from which vulnerability databases can be downloaded, adds support for parsing Dockerfile heredoc syntaxes, and introduces schema enhancements for container and IaC scan results.

— Lists every change but no single invocation example. CLI updates (2026-10-27)

07 **`--include` module selection in Pipeline Scan** NEW

75

Pipeline Scan adds an `--include` parameter to specify top-level modules to include during scanning, with prescan module discovery now shown in results; available in Veracode CLI version 23.4.3-0 and Docker image version 23.4.3.

— Names the exact flag and versions, no sample command. SAST updates archive (2026-10-27)

08 **Azure DevOps Upload and Scan task parameters** NEW

75

The Veracode Azure DevOps Extension's Upload and Scan task gains two optional parameters: `scanpollinginterval` to control how often the task polls for scan status, and `lifecyclestage` (not supported in YAML pipelines).

Control polling frequency for a long-running Upload and Scan in the Azure DevOps pipeline task to reduce API call overhead.

```yaml
- task: Veracode@3
  inputs:
    ConnectionDetailsSelection: 'Endpoint'
    AnalysisService: 'VeracodeEndpoint'
    veracodeAppProfile: 'MyApp'
    version: '$(build.buildNumber)'
    filepath: 'drop/myapp.zip'
    scanpollinginterval: '60'
```

— Has exact flag names and a working YAML example.

Integration updates archive (2026-10-27)

09 **Fire-and-Forget mode in `veracode fix sca`** NEW 60

`veracode fix sca` supports asynchronous Fire-and-Forget (FnF) mode, submitting remediation requests for background processing without waiting for fixes or pull requests.

— Names the command and mode, no example invocation. CLI updates (2026-10-27)

10 **Python CWE-201/CWE-918 recategorization** BREAKING 60

Python applications with flaws previously reported as `CWE-201` that should have been `CWE-918` are being recategorized, potentially breaking existing flaw matching and requiring new mitigations.

— Names exact CWE codes and consequence, no remediation steps.

SAST updates archive (2026-10-27)

11 **Unstable status option in Jenkins Plugin** NEW 60

The Veracode Jenkins Plugin adds a Show Unstable Status for Failed Policy Evaluation option that marks a job as Unstable when a scan completes but fails the security policy.

— Describes exact option name and behavior, no setup steps.

Integration updates archive (2026-10-27)

THINNER COVERAGE BELOW

12 **New COBOL scanner for Static Analysis** NEW 55

Veracode introduces a new COBOL scanner for Static Analysis using advanced pattern recognition and static analysis techniques for more accurate vulnerability detection in COBOL code; all COBOL scans now use this upgraded scanner, which may surface additional vulnerabilities and impact flaw matching for existing applications.

— Explains mechanism and impact but no migration guidance.

SAST updates archive (2026-10-27)

13 **New language and framework support in Static Analysis** NEW

55

Static Analysis adds support for Quarkus (Kubernetes-native Java stack for OpenJDK HotSpot and GraalVM), Micronaut 3.8.x, `.NET Minimal API`, PHP Laravel 10, initial `iOS 17` support, `Go 1.21`, AWS SDK for Go v2, and `KMS` support for AWS SDK for Go.

— Lists every named framework/language but no scan guidance.

SAST updates archive (2026-10-27)

---

14 **Azure DevOps Flaw Import task enhancements** IMPROVED

55

The Flaw Importer task in Veracode Azure DevOps Extension 3.23.0 adds an option to limit imported findings to Static Analysis and SCA only, an option to overwrite the iteration path on work items of imported flaws, and a new iteration path field in the Flaw Importer task's work item settings.

— Names the task and version but no flag identifiers given.

Integration updates archive (2026-10-27)

---

15 **GitHub Workflow App for org-wide scanning** NEW

55

Veracode releases a GitHub Workflow App enabling Static Analysis, SCA, and Container Security scans across thousands of GitHub repositories from a single centralized workflow location.

— Explains scope and scan types but no setup instructions.

Integration updates archive (2026-10-27)

---

16 **Target UUIDs for `veracode dynamic scan`** IMPROVED

50

The `veracode dynamic scan` command now supports target UUIDs for triggering scans, while webhook UUIDs remain supported.

— Names the command and option but no example. CLI updates (2026-10-27)

---

17 **New Dynamic Analysis scan statuses** NEW SUPPORTED PLATFORM NEW

50

Adds Dynamic Analysis scan statuses `STOPPED_VERIFYING_PARTIAL_RESULTS` and `STOPPED_PARTIAL_RESULTS_AVAILABLE`, surfaced through the Java API Wrapper and Jenkins Plugin.

— Names both statuses but no guidance on handling them.

Integration updates archive (2026-10-27)

---

18 **SCA plugin for JetBrains IDEs** NEW

50

Veracode SCA Scan for JetBrains 0.6.0 is a new extension integrating Software Composition Analysis into IntelliJ IDEA and PyCharm to detect vulnerabilities in open-source libraries, dependencies, and licenses.

— Names IDEs and version but no install/usage path.

Integration updates archive (2026-10-27)

---

19 **TSRV mitigations and .NET 7 in Visual Studio plugin** IMPROVED

45

Veracode Static for Visual Studio 2019 and 2022 v1.9.0 adds support for TSRV mitigations and .NET 7.

— Names versions and targets but no further mechanism.

Integration updates archive (2026-10-27)

---

20 **TFS support ends in Azure DevOps Extension** BREAKING

40

Veracode ends support for Team Foundation Server (TFS) in the Azure DevOps Extension as of version 3.24.0; pipelines relying on TFS integration will need to migrate.

— States the change and version but no migration path detail.

Integration updates archive (2026-10-27)

UPDATED DETECTION Scan engine false positive reduction across CWE categories

NEW DETECTION Improved detection for CAA record checks on SSL certificates for CWE-862

NEW DETECTION New detections for open ports, HTTP headers, TLS, and fingerprinting

NEW INTEGRATION EASM integration with DAST for automatic target creation

NEW INTEGRATION Postman Collections support for API scanning

UPDATED SUPPORTED PLATFORM ISM endpoint 23.12.1 adds Java 21 support and virtual threading

└──▷ BREAKING ON UPGRADE

! End of support for Team Foundation Server (TFS) in Veracode Azure DevOps Extension 3.24.0.

! Veracode deprecated support for .NET cleansers for flaw classes CWE-93, CWE-113, and CWE-117 across these functions: `antixsslibrary.dll: Microsoft.Security.Application.AntiXss.HtmlAttributeEncode`, `antixsslibrary.dll: Microsoft.Security.Application.AntiXssEncoder.HtmlAttributeEncode`, `antixsslibrary.dll: Microsoft.Security.Application.Encoder.HtmlAttributeEncode`, `antixsslibrary.dll: Microsoft.Security.Application.Encoder.HtmlEncode`, `mscorlib.dll: System.Security.SecurityElement.Escape`, `system.dll: System.Net.WebUtility.HtmlEncode`, `system.web.dll: System.Web.HttpServerUtility.HtmlEncode`, `system.web.dll:`

`System.Web.Security.AntiXss.AntiXssEncoder.HtmlEncode`, `system.web.dll:` `System.Web.Util.HttpEncoder.HtmlAttributeEncode`, `system.web.dll:` `System.Web.Util.HttpEncoder.HtmlEncode`, `system.web.mvc.dll:` `System.Web.Mvc.HtmlHelper.AttributeEncode`, `system.web.mvc.dll:` `System.Web.Mvc.HtmlHelper.Encode`, `system.windows.browser.dll:` `System.Windows.Browser.HttpUtility.HtmlEncode`, `system.windows.dll:` `System.Net.HttpUtility.HtmlEncode`, `System.Runtime.dll:` `System.Net.WebUtility.HtmlEncode` — applications relying on these as cleansers will no longer have those flaws suppressed.

! Python applications with flaws previously reported as `CWE-201` that should have been `CWE-918` will be recategorized, potentially breaking existing flaw matching and requiring new mitigations.

! All COBOL scans now use the upgraded scanner, which may identify additional vulnerabilities and impact flaw matching for existing applications.

`DAST updates (2026-10-27)` **BLOG ↗** **STABLE**

Veracode DAST gains domain-based HTTP basic auth, custom host-to-IP mappings, multi-team target ownership, Directory Restrictions for scan scoping, advanced scheduling, and a public REST API for target management.

Adds an optional `Domain` field when configuring HTTP basic authentication for DAST web application scans, enabling authentication to applications that require domain-based credentials.

Adds support for custom host-to-IP address mappings in DAST scans, configurable when DNS lookup is unavailable or not desired.

Enables assigning multiple teams to a DAST target when creating a web application or API specification, improving collaborative security ownership.

Adds `Directory Restrictions` to DAST target configuration to control scan scope.

Adds enhanced scheduling capabilities to the DAST `Schedule` feature, with advanced configuration options and greater flexibility for managing scan schedules.

`+14 more`

└─▷ **BREAKING ON UPGRADE**

! The DAST scheduling REST API endpoints have changed. Integrations using the REST API to manage DAST schedules must be updated to use the new endpoints as defined in the DAST API specification.

WAS THIS USEFUL? 👍 👎

Phoenix Security

Feedback ⓘ

Phoenix DAST Scanner – Configuration Guide (2026-10-27) BLOG ↗ RANK

Phoenix Security correlates scanner, cloud and runtime findings with exploitability and ownership, and uses AI agents to prioritise and open remediation pull requests.

> Phoenix added a native DAST Scanner integration for scheduling and monitoring authenticated web application scans directly from the platform.

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

01 **Phoenix DAST Scanner integration** NEW 82

A new DAST Scanner integration appears under `Integrations > Scanners > WEB`, letting teams schedule authenticated web scans and manage targets from within Phoenix. Each target is configured with a target URL, start time, time zone and schedule, with multiple targets added via `+ Add New Target`, and authentication via the Auth Token (Bearer) field supports bearer and JWT tokens (cookie, basic, and client-certificate auth are not supported). The integration page monitors target status — Active/Fetched, In Progress, Not Started, Disabled, Missing, and Error — alongside last fetch date, and exposes Fetch and `Fetch/pause all` controls to enable or pause individual or all targets, plus an optional setting to auto-delete missing targets not reported for a configurable number of days.

Set up a scheduled authenticated DAST scan for a web target that requires a bearer token for login.

> 📍 In Phoenix, go to Integrations › Scanners › WEB › Phoenix DAST Scanner. Expand Target 1, enter the target URL, set Start Time and Schedule, and paste the bearer or JWT token into the Auth Token (Bearer) field. Select Create Scanner.

Pause all DAST target processing during a maintenance window without removing the integration.

> 📍 Open the created scanner integration page in Phoenix. Select Fetch/pause all to pause processing for every configured target. Select it again to re-enable when the window closes.

– Names exact menu path, field names, states, and controls.

Phoenix DAST Scanner – Configuration Gu…

WAS THIS USEFUL?  

Black Duck

Feedback ⓘ

Introducing Signal Agentic Fix: Your 24×7 remediation team (2026-10-27) BLOG ↗

RANK

Black Duck helps organizations secure their software supply chain by providing deep visibility into open source components, licenses, and vulnerabilities.

> Black Duck launched Signal Agentic Fix, an AI remediation agent that autonomously scans, triages, patches, validates, and submits pull requests for vulnerabilities, alongside a companion mode that plugs into coding agents like Claude Code and GitHub Copilot.

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

01 Signal Agentic Fix autonomous remediation agent NEW

72

Signal Agentic Fix is an AI agent that performs end-to-end vulnerability remediation—scan, triage, fix generation, validation, and pull-request submission—without human bottlenecks. It reuses the same deterministic SAST and SCA engines from Polaris, Coverity, and Black Duck SCA for consistent scan results inside the agent workflow, and is available in early access as a Claude Code plug-in for Coverity users, with additional coding-agent and DevOps pipeline support planned for Coverity, Polaris, and Black Duck SCA. A planned addition, Signal AI Code Analysis support, will enable deeper discovery of novel risks such as IDOR and BOLA within agent runs.

– Names engines and rollout path but mechanism of fix generation is vague

Introducing Signal Agentic Fix: Your 24…

THINNER COVERAGE BELOW

02 Agentic Fix for Coding Agents companion NEW

45

Agentic Fix for Coding Agents integrates as a companion to tools like Claude Code and GitHub Copilot, letting developers remove defects in existing code and prevent new ones as AI writes code.

– States purpose and named integrations but no setup or usage detail

Introducing Signal Agentic Fix: Your 24…

WAS THIS USEFUL? 👍 👎

GitLab Feedback ⓘ

Extend GitLab Duo Agent Platform into your business systems (2026-10-27) BLOG ↗

RANK

Teams and agents ship trusted software faster with GitLab, the intelligent orchestration platform with context and governance for the entire software lifecycle.

> GitLab shipped a GA Flows API that lets any external system trigger custom Duo Agent Platform flows to auto-triage issues, alongside the YAML-based flow definition format and new authoring surfaces in VS Code and the project UI.

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

01 **Flows API to trigger custom AI agent flows** NEW 95

The `POST /api/v4/ai/duo_workflows/workflows` endpoint (generally available from GitLab 19.5) lets any external system trigger a custom Duo Agent Platform flow over HTTP by passing `project_id`, `ai_catalog_item_consumer_id`, and a `goal`. Request body fields `allow_agent_to_request_user`, `agent_privileges`, and `pre_approved_agent_privileges` control whether the flow executes autonomously and which tool permissions it is granted. The response includes a `web_url` field linking to the agent session for inspecting each step the flow took, and the flow's `ai_catalog_item_consumer_id` can be looked up via a GraphQL query.

Trigger a custom Duo Agent flow from any external system to auto-triage an error and open a GitLab issue — no human required.

```
$ curl --request POST
https://gitlab.com/api/v4/ai/duo_workflows/workflows \
  --header 'PRIVATE-TOKEN: <your_project_access_token>' \
  --header 'Content-Type: application/json' \
  --data '{
    "project_id": 12345,
    "ai_catalog_item_consumer_id": 67,
    "goal": "Class: OpportunityDiscountService\nException:
System.NullPointerException\nMessage: Attempt to de-reference a
null object\nStack trace: ...\n",
    "start_workflow": true,
    "allow_agent_to_request_user": false,
    "agent_privileges": [1, 2, 3],
    "pre_approved_agent_privileges": [1, 2, 3]
  }'
```

Enable a custom flow and retrieve its `ai_catalog_item_consumer_id` for use in API calls, then navigate to AI > Flows in your project.

> 📍 1. In your GitLab project, go to AI › Flows. 2. Open the `Managed` tab and select your flow. 3. Select `Enable` and choose the project that holds your source code. 4. Run the GraphQL query from the docs to retrieve the gid:/`/gitlab`/AiCatalogItemConsumer/<number> and note the numeric ID as your ai_catalog_item_consumer_id. 5. Create a project access token with the `api` scope and at least the Developer role for use in the PRIVATE-TOKEN header.

— Full endpoint, fields and response documented with a runnable curl example

Extend GitLab Duo Agent Platform into y…

02 **Custom flow definitions in flow registry v1 YAML** NEW

73

Custom agent flows are defined as a single YAML file using `flow registry v1` syntax, specifying `AgentComponent` agents, toolsets (`find_files` , `grep` , `read_file` , `list_commits` , `create_issue`), and prompt templates that drive the agent's behavior.

Define a custom AI triage agent flow that reads repository files and creates a GitLab issue, using flow registry v1 YAML syntax.

```yaml
version: "v1"
environment:
  ambient
components:
  - name: "sf_error_triage"
    type: AgentComponent
    prompt_id: "sf_error_triage_prompt"
    inputs:
      - from: "context:project_id"
        as: "project_id"
      - from: "context:goal"
        as: "error_report"
    toolset:
      - "find_files"
      - "grep"
      - "read_file"
      - "list_commits"
      - "create_issue"
prompts:
  - prompt_id: "sf_error_triage_prompt"
    prompt_template:
      system: |
        You triage errors for a team. Find the source file, read
the code,
        list recent commits, and create one issue with root-cause
analysis.
      user: |
        Project ID: {{project_id}}
        Error report: {{error_report}}
      placeholder: history
routers:
  - from: "sf_error_triage"
    to: "end"
flow:
  entry_point: "sf_error_triage"
```

— Names exact YAML syntax and toolset but no storage path given

Extend GitLab Duo Agent Platform into y…

THINNER COVERAGE BELOW

03 **Flow authoring in VS Code and project UI** NEW 45

Custom flows can be created and tested in VS Code using the GitLab flow builder, and managed in-project via the new AI > Flows > New flow UI, which includes an Enable action under the Managed tab to scope a flow to a specific project.

— Names UI path and tool but gives no deeper mechanism

Extend GitLab Duo Agent Platform into y…

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ announced · Oct 26, 2026

BLOG ↗ RANK

Sonarqube is a enforce consistent quality gates across every pull request and pipeline, from IDE to CI/CD, with automated code review for bugs and security issues.

> SonarQube's Remediation Agent now extends to Hunter Agent findings, letting teams filter, auto-fix, and CI-validate fixes for broken access control, business logic flaws, and auth gaps across both SonarQube Cloud and self-deployed SonarQube Server.

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

01 **Remediation Agent support for Hunter Agent findings** NEW

88

SonarQube Cloud's Issues page now supports the `hunter-agent` tag and category tags like `hunter-broken-access-control` to filter Hunter Agent findings (broken access control, business logic flaws, auth gaps) alongside the existing Fixable by Agent filter; up to 20 findings can be selected and sent to Fix with Agent. Each Hunter Agent finding generates its own individual pull request (or merge request on GitHub), unlike the grouped-by-rule-key behavior used for rule-based issues. For GitHub Actions pipelines, CI follow-up is enabled on these fix PRs: when CI fails, the agent reads build logs, pushes a corrective commit, and retries up to three times before leaving the PR open. This capability is also brought to SonarQube Server 2026.5, with both the Hunter Agent and Remediation Agent able to run on self-deployed infrastructure.

Filter the Issues page to only Hunter Agent authorization findings ready for automated remediation, then dispatch them for fixing.

> 📍 In SonarQube Cloud, open your project Issues page, select the main branch, enable the Fixable by Agent filter, then add the tag `hunter-broken-access-control` (or `hunter-agent` for all Hunter Agent findings). Select up to 20 findings and choose Fix with Agent.

— Names exact tags, filter, PR behavior, retry limit, and server version

launch-20261026-7b057b78

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

kingfisher

Feedback ⓘ v2.15.0 NOTES ↗ RANK

Detect secrets + live validation. Map blast radius.

> Kingfisher expands its blast-radius analysis with multi-region/multi-project scope flags, an automatic capability-based risk and coverage assessment in report output, and a new HTML investigation viewer for triaging findings.

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

01 Multi-region and multi-project scope expansion for blast-radius scans

NEW 90

Adds `--aws-region` to the `scan` and standalone `toon` commands to expand AWS discovery beyond the SDK-selected region, accepting up to eight explicit regions (global IAM/S3 resources are inventoried once). Adds `--gcp-project` to `scan` and `toon` to expand GCP discovery beyond the home project, accepting up to seven additional projects, each using its own visible IAM context.

Expand AWS blast-radius discovery across multiple regions when a leaked key may have cross-region resources.

```
$ kingfisher scan --aws-region us-east-1 --aws-region eu-west-1 --aws-region ap-southeast-1 ./key.json
```

Expand GCP blast-radius discovery to additional projects beyond the credential's home project to catch cross-project role bindings.

```
$ kingfisher toon ./key.json --gcp-project production-project --gcp-project staging-project
```

— Exact flags, numeric limits, and runnable commands given product docs

02 Capability-based risk and coverage assessment in blast-radius reports

NEW 85

Blast-radius output now includes automatic risk and coverage assessment in `scan` and `toon` (no extra flag needed), labeling permissions by capability category — e.g. `s3:ListBucket` as Metadata read, `s3:GetObject` as Sensitive data read — and distinguishing direct access, candidate role paths, and observed API access from mere

policy grants or conditions. This is exposed via the new `provider_metadata.authorization_evidence.assessment` field, containing schema version, assessment timestamp, capability categories, evidence labels, coverage, resource explanations, recommendations, and notes; `capability_risk` is tracked separately from the original `severity` field. Detailed assessments are capped at 4096 capability records and 1000 resources per identity, with up to 256 permissions per resource, and legacy report data is retained separately. The `observed_at` timestamp is intentionally excluded from snapshot comparisons since it varies between runs.

— Field name and caps given but no query example shown `product docs`

---

03 **HTML investigation viewer for blast-radius reports** NEW

80

Adds a blast-radius HTML investigation viewer with Overview, Findings, Coverage, and Blast Radius views that opens existing JSON/JSONL reports and supports drag-and-drop via a compact import bar. It groups identities by provider and surfaces policy provenance, candidate paths, conditions, exclusions, and observed API methods alongside a separate resource map. A synthetic investigation report illustrating AWS, GCP, Azure, Hugging Face, and GitHub evidence can be opened with `kingfisher view` or the report picker to explore the viewer before running a real scan.

Open the built-in synthetic investigation report to explore the HTML triage viewer before running a real scan.

```
$ kingfisher view
```

— Command to launch viewer given; UI detail stays descriptive `product docs`

NEW INTEGRATION Betterleaks v2 program support

v2.15.0 NOTES ↗ LATEST STABLE

Kingfisher v2.15.0 adds PDF scanning, `jwt_sign` for RS256 JWT construction, `--fallbacks` for HTTP fallback environments, versioned blast-radius assessments, provider-issued repo IDs in Git metadata, and `KF_GITHUB_MAX_RETRY_WAIT_SECS`.

└──▷ GET THIS VERSION ▾

Homebrew

```
$ brew install kingfisher
# already installed? upgrade to this version:
$ brew upgrade kingfisher
```

Docker

```
$ docker pull ghcr.io/mongodb/kingfisher:2.15.0
```

git

```
$ git clone --branch v2.15.0
https://github.com/mongodb/kingfisher.git
# already have the repo? check out this version:
$ git checkout v2.15.0
```

Homebrew lists this version · published under the project's GitHub account · checked 28 Oct

Wrong command?

Adds `KF_GITHUB_MAX_RETRY_WAIT_SECS` environment variable to configure how long scans wait through GitHub quota resets.

Adds `--fallbacks` CLI flag, `http_fallbacks=True` Python option, and .http_fallbacks(true) Rust option to enable up to two opt-in HTTP fallback environments for validation and revocation.

Adds `--max-file-size` enforcement (default 256 MB) on Git-history blobs before reading content; disable with `--max-file-size 0` or `--no-limits`.

Adds `jwt_sign` function to construct complete RS256 JWTs from JSON claims and PEM keys.

Adds offline, resource-bounded PDF page-text extraction for CLI file, Git-history, and archive scans; Python users opt in with expand_content(..., pdf=True).

+3 more

└─▷ **BREAKING ON UPGRADE**

! Legacy HTTP `request.url: [primary, secondary, ...]` lists now use only the primary URL by default; enabling `--fallbacks` tries at most the primary and two alternatives.

! HTTP revocation no longer follows redirects, including for legacy direct-primary rules; Anthropic, Terraform, and DigitalOcean access-map requests also refuse redirects — configure final canonical URLs instead.

! Rust callers of rules/scanner crates must follow the 3.0 migration guide (`docs/LIBRARY.md#migrating-to-rules-and-scanner-30`) for public schema changes.

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS

Feedback ⓘ

Scan developer machines with less noise, and clean up what you find (2026-10-27)

BLOG ↗ RANK

GitGuardian SaaS is a cloud-hosted security platform that detects exposed secrets, API keys, and credentials in source code repositories and developer workflows.

> GitGuardian SaaS's machine protection now supports low-priority background scanning on developer machines, adds an Overview dashboard to guide remediation, redacts more AI-agent history, expands validity checks to 501 secret types, and cuts noise from IDE history and build/cache folders — alongside a breaking change to how custom hosts are configured for validity checks.

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

01 Background-priority scanning for machine protection NEW

95

Adds a `--background` flag to `ggshield machine scan` and `ggshield machine report` to run scans at low priority using 2–4 threads and one-tenth of normal memory (minimum 512 MB), suited for scheduled MDM jobs. Adds `scan.resources.background.mode: auto` to `machine_scan.yaml` to make background-priority scanning the default for every scan regardless of how it's triggered.

Run a scheduled daily inventory scan at low priority so developer machines stay responsive during the MDM job.

```
$ ggshield machine report --background
```

Make background-priority scanning the permanent default so every scan — however triggered — leaves the machine responsive.

```yaml
scan:
  resources:
    background:
      mode: auto
```

— Names flags, config key and concrete resource limits, with runnable examples

Scan developer machines with less noise…

02 **Overview page in machine dashboard** NEW 65

`ggshield machine dashboard` gains an Overview page that lists one row per cleanup area (AI agents, repositories, history and logs) and guides remediation order: overprivileged secrets first, then history leaks, then secrets in daily use.

Open the local remediation dashboard to see the Overview and start cleaning up findings grouped by AI agent or repository.

```
$ ggshield machine dashboard
```

— Names the CLI entry point and remediation ordering but no deeper mechanism

Scan developer machines with less noise…

THINNER COVERAGE BELOW

03 **Validity checks expanded to 501 secret types** IMPROVED 55

Validity checks now cover 501 secret types, up from 432, adding coverage for Elastic, Datadog, Langfuse, Sumo Logic, Gitea, and Azure Container Registry.

— Precise numbers and named providers but nothing for a reader to do

Scan developer machines with less noise…

04 **Noise reduction for IDE editor-history copies** IMPROVED 55

Findings from VS Code, Cursor, VSCodium, Windsurf, and JetBrains IDE editor-history copies are now suppressed unless the secret in them is valid, reducing duplicate noise.

— Clear mechanism and named editors but passive, no user action

Scan developer machines with less noise…

05 **Custom hosts for validity checks now workspace-only** BREAKING 55

Validity checks now take custom hosts only from your GitGuardian workspace; environment variables and local config files that previously set custom hosts are no longer read.

— Clear before/after but no migration command given

Scan developer machines with less noise…

06 **Expanded folder exclusions for build and cache leftovers**

IMPROVED 50

Scanning now skips build leftovers (`pip-*` , `pytest-of-*` , `_MEI*`), Windows package caches (NuGet, Scoop), and Electron app caches, while still scanning

credential-holding folders.

— Names exact patterns excluded but offers no configurable surface

Scan developer machines with less noise…

07 **Codex and opencode history redaction** NEW 45

The local dashboard can now redact Codex and opencode history, including their SQLite history databases.

— States the capability but gives no command or workflow to trigger it

Scan developer machines with less noise…

└──▷ **BREAKING ON UPGRADE**

! Validity checks now take custom hosts only from your GitGuardian workspace; environment variables and local config files that previously set custom hosts are no longer read.

WAS THIS USEFUL?  

◆ Web Application Security

WPScan

Feedback ⓘ v4.2.0 NOTES ↗ CODE ↗ RANK

WPScan WordPress security scanner. Written for security professionals and blog maintainers to test the security of their WordPress websites. Contact us via contact@wpscan.com

> WPScan v4.2.0 adds a configurable cap on response read size to guard against memory exhaustion, and improves SARIF output to correctly flag aborted scans as failed runs.

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

01 Response size cap via --max-response-size

BREAKING

93

The new `--max-response-size MiB` flag stops reading any HTTP response once it exceeds the specified size, guarding against memory exhaustion on targets serving huge or endless responses; default is 16 MiB, max is 2047, and `0` disables the cap entirely. This changes default behavior: responses larger than 16 MiB are now truncated by default, so existing scans relying on unbounded reads must pass `--max-response-size 0` to restore the previous behavior. The option can also be persisted via `max_response_size` under `cli_options` in a project config file such as `pwd/.wpscan/scan.yml`.

Scan a target that serves unusually large pages without risking memory exhaustion — raise the cap to 64 MiB instead of the 16 MiB default.

```
$ wpscan --url https://target.tld --max-response-size 64
```

Persist a raised response-size cap in a project config file so every scan in this directory inherits it without repeating the flag.

```yaml
cli_options:
  url: 'https://target.tld'
  max_response_size: 64
```

— Names exact flag, defaults, limits, config key and migration path v4.2.0

THINNER COVERAGE BELOW

02 SARIF aborted-run reporting

IMPROVED

55

SARIF output now reports aborted runs as unsuccessful, recording the abort reason in `toolExecutionNotifications`, so downstream consumers show an aborted scan as a failed analysis rather than a clean run with zero alerts.

– Explains mechanism and SARIF field but no CLI usage shown v4.2.0

└──▷ **BREAKING ON UPGRADE**

! Responses larger than 16 MiB are now truncated by default; use `--max-response-size 0` to restore the previous unbounded behavior.

WAS THIS USEFUL? 👍 👎

PortSwigger Burp Suite Feedback ⓘ

DAST 2026.10.2 (2026-10-27) BLOG ↗ RANK

Burp suite is a we enable the world to secure the web. Burp Suite, the Web Security Academy, and world-leading research from PortSwigger â the AppSec communityâs home.

> Burp Suite DAST 2026.10.2 adds ServiceNow ticketing integration, expands API scan coverage, and reduces scan memory usage for better reliability.

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

01 **ServiceNow ticketing integration** NEW 40

Burp Suite DAST now integrates with ServiceNow, enabling security issues found by Burp Suite DAST to be raised directly in ServiceNow tickets.

– Names the integration but no setup or config detail DAST 2026.10.2 (2026-10-27)

02 **Expanded API scan coverage** IMPROVED 20

Burp Suite DAST expands API scan coverage to detect more issues across API targets.

– No specifics on which issue types or APIs DAST 2026.10.2 (2026-10-27)

03 **Reduced scan memory usage** IMPROVED 20

Burp Suite DAST reduces memory usage during scans to improve scan reliability.

– No numbers or mechanism given for the reduction DAST 2026.10.2 (2026-10-27)

WAS THIS USEFUL? 👍 👎

caido

Feedback ⓘ announced · Oct 26, 2026 BLOG ↗ RANK

Caido releases, wiki and roadmap

> Caido's v0.59.0 release centers on a new AI provider framework supporting major LLM vendors and local custom endpoints, alongside an AI-assisted Workflow generator, an experimental traffic capture/passthrough system, and quality-of-life improvements to Replay, Desktop multi-instancing, and the macOS menu bar app.

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

01 AI provider configuration and credential testing NEW

82

Adds `Settings > AI > Providers` to configure OpenAI, Anthropic, Google Gemini, SpaceXAI, OpenRouter, and Amazon Bedrock from one place with credential testing before use. Also supports custom providers (e.g. a local LM Studio endpoint) by setting a Provider ID, matching API type, Base URL, and an optional API Key, with a Test action to verify the connection.

Connect a local LM Studio endpoint as a custom AI provider without an API key, so your prompts never leave your machine.

> 📍 In Caido, go to Settings › AI › Providers › Custom provider. Set Provider ID to a label of your choice, select the matching API type, enter your LM Studio Base URL (e.g. `http://localhost:1234/v1`), and leave API Key empty. Click Test to verify the connection.

— Names exact settings path, six providers, and custom-provider fields with a worked example.

launch-20261026-69a40eac

02 Capture modes and passthrough in Intercept NEW

80

Adds an `Intercept > Capture` menu (marked Experimental) with three recording modes — All traffic, In-scope only, and Custom rules — plus a Passthrough list under Custom rules mode so hosts on it pass through Caido without being recorded. In-scope only uses the active Scope definition to silently skip recording out-of-scope requests.

Keep noise out of a project by recording only in-scope hosts during an engagement.

> 📍 In Caido, open Intercept › Capture and select In-scope only. Caido will use the active Scope definition and silently pass through any out-of-scope requests without recording them.

— Exact menu path, three named modes, and a passthrough mechanism with steps.

launch-20261026-69a40eac

THINNER COVERAGE BELOW

03 **Generate action for Workflows** NEW 50

Adds a Generate action in the Workflows page to produce an editable passive, active, or convert Workflow from a plain-language description.

.

— Clear UI entry point but no worked example or deeper mechanism given.

launch-20261026-69a40eac

04 **macOS menu bar app reconnect behavior** IMPROVED 40

Reworks the macOS menu bar app to reopen the Instance Manager and reconnect to a running local instance after its window is closed.

— States the behavior change but no steps to try it. launch-20261026-69a40eac

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

DataDome

Feedback ⓘ Cloudflare Worker Rust module 1.0.0 (2026-10-27)

BLOG ↗ RANK

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

> DataDome's Fraud SDK .Net 2.1.0 adds an optional field to let callers specify the effective time of a fraud event independent of submission time.

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

01 **eventDate field in Fraud SDK .Net payload** NEW 60

The Fraud SDK .Net 2.1.0 adds an optional `eventDate` field to the event payload to specify the effective event time, useful for replaying or backfilling events rather than relying on submission time.

Pass `eventDate` in a fraud event payload to record the actual time an event occurred rather than the submission time — useful for replaying or backfilling events.

```json
{
  "eventDate": "2026-10-27T14:32:00Z",
  "...": "..."
}
```

— Names the exact field and payload usage with example. Fraud SDK .Net 2.1.0 (2026-10-27)

NEW INTEGRATION DataDome Cloudflare Workers integration (Rust)

> Cloudflare Worker Rust module 1.0.0 (2026-10-27) BLOG ↗ LATEST STABLE
>
> Adds the `datadome-cloudflare-worker` Rust crate for native DataDome bot protection inside Cloudflare Workers.

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ Fastly Terraform Provider (beta) 0.2.6 (2026-10-27)

BLOG ↗ RANK

Fastly is an edge cloud platform offering content delivery network, load balancing, edge computing, and web application security services for developers and enterprises.

> Fastly's biggest change this window is automated virtual patching for Next-Gen WAF workspaces, surfaced both as a workspace setting and in the Terraform provider, alongside a Terraform migration path off legacy VCL services and several Rust SDK additions covering client detection, product enumeration, shield timeouts, and Image Optimizer fallback behavior.

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

01 **VPN/proxy client detection in Rust SDK** NEW 75

Adds `Request::get_vpn_proxy_data`, `Request::get_vpn_proxy_*`, and `RequestHandle::vpn_proxy_*` methods to expose metadata for detecting whether a downstream client is residential, a VPN, or a proxy.

— Exact method names given, no sample usage code Rust SDK 0.13.2 (2026-10-27)

02 **Image Optimizer fallback behavior on disabled services**

IMPROVED 65

When `Request::with_image_optimizer` is used on a service where Image Optimizer is not enabled, the origin response is now passed through unmodified with a `Fastly-Io-Error` response header instead of failing.

— Clear before/after behavior with named header Rust SDK 0.13.2 (2026-10-27)

THINNER COVERAGE BELOW

03 **Product enumeration via compute_runtime::products** NEW 55

Adds `fastly::compute_runtime::products` to report which products are enabled on the service.

— Named function but no detail on return shape or usage Rust SDK 0.13.2 (2026-10-27)

04 **VCL-to-CDN-auto state migration path** NEW 50

Adds a proof-of-concept state migration path from legacy `fastly_service_vcl` resources to `service_cdn_auto` in the Fastly Terraform Provider (beta).

— Names both resource types but no migration steps or caveats

Fastly Terraform Provider (beta) 0.2.6 …

05 **Between-bytes timeout for shield backends** NEW 35

The Rust SDK now supports adding a between-bytes timeout for shield backends.

— No config key, method name, or default value given `Rust SDK 0.13.2 (2026-10-27)`

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Yamato Security Hayabusa

Feedback ⓘ v4.3.0 NOTES ↗

CODE ↗ RANK

Hayabusa (隼) is a sigma-based threat hunting and fast forensics timeline generator for Windows event logs.

> Hayabusa v4.3.0 adds a fixed UTC-offset timestamp option across six commands and extends unaliased rule-field resolution to additional event log fields.

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

01 Fixed UTC offset timestamps via --timezone flag NEW

90

The `-Z, --timezone <OFFSET>` flag was added to `dfir-timeline`, `search`, `eid-metrics`, `logon-summary`, `log-metrics`, and `extract-base64` to render output timestamps with a fixed UTC offset in signed `+HH:MM` or `-HH:MM` form (e.g. `--timezone "+09:00"` or `-Z "-01:00"`). It conflicts with `--utc` and `--iso-8601`, and filtering, sorting, and histogram bins remain UTC-based.

— Exact flag, affected commands, and example values given v4.3.0

THINNER COVERAGE BELOW

02 Unaliased rule-field resolution for EventData and UserData

IMPROVED

30

Rule-field resolution without aliases now extends to the `EventData` and `UserData` fields.

— Named but no mechanism or example given v4.3.0

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-10-26 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 extracts multi-tenancy into a self-contained plugin library and adds a lightweight plugin hook registry for MongoDB, SQL, and Elasticsearch filters.

└──▷ USE IT

Register a custom MongoDB query filter plugin to enforce tenant-scoped result isolation without patching CAPEv2 core.

```python
from lib.cuckoo.common.hooks import register_mongo_filter

def my_tenant_filter(query: dict) -> dict:
    query['info.tenant_id'] = MY_TENANT_ID
    return query

register_mongo_filter(my_tenant_filter)
```

Register a SQL task filter to restrict list_tasks() results to a specific tenant, replacing the removed inline `tenant_id` filtering.

```python
from lib.cuckoo.common.hooks import register_sql_task_filter

def my_sql_filter(query):
    from lib.cuckoo.core.data.task import Task
    return query.filter(Task.owner == MY_USER_ID)

register_sql_task_filter(my_sql_filter)
```

Adds `lib/cuckoo/common/hooks.py` : a lightweight plugin hook registry exposing `register_mongo_filter` , `register_sql_task_filter` , `register_es_filter` , and `register_report_hook` , replacing inline tenancy logic with zero-overhead hooks that pass queries through unchanged when no plugin is registered.

Integrates `lib/tenancy/` as a self-contained library for multi-tenancy and multi-cloud object storage, extracted from CAPEv2 core into the standalone `CAPESandbox/CAPE-tenancy` plugin repository.

Adds Django Group-based sandbox resource entitlement policy in `lib/tenancy/resources.py` , controlling which VMs, routes, tags, interactive sessions,

and timeouts a user may discover in the UI/API and request at task submission time.

Wires `apply_mongo_filter` into `dev_utils/mongodb.py` so registered MongoDB query filters are automatically applied in `mongo_find` and `mongo_find_one` (bypass with `no_hooks=True`).

Wires `apply_sql_task_filter` into `lib/cuckoo/core/data/tasking.py` so registered SQL task filters are applied in list_tasks() and count_matching_tasks().

+1 more

└─▷ **BREAKING ON UPGRADE**

- ! The `[multitenancy]`, `[central_mode]`, `[centralstore]`, and `machine_info_token` stanzas are removed from `conf/default/cuckoo.conf.default`, `conf/default/api.conf.default`, and `conf/default/reporting.conf.default`; deployments relying on these config stanzas must migrate to the `CAPESandbox/CAPE-tenancy` plugin.
- ! `Task.tenant_id` and `Task.visibility` columns are removed from `lib/cuckoo/core/data/task.py`; code or queries referencing these fields will break (the Alembic migration `3a1b_tenant_visibility` is preserved as a no-op for continuity).
- ! The `lock_engine` attribute is removed from `lib/cuckoo/core/database.py`; any custom code referencing it will break.
- ! The standalone web modules `web/web/tenancy_optional.py` and `web/users/tenancy.py` are deleted; imports of these modules will fail.

WAS THIS USEFUL? 👍 👎

VirusTotal YARA

Feedback ⓘ v4.6.0 NOTES ↗ CODE ↗

The pattern matching swiss knife

Adds `yr_get_version` C API to query the runtime libyara version before `yr_initialize`, enabling shared-library version checks.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v4.6.0 https://github.com/VirusTotal/yara.git
# already have the repo? check out this version:
$ git checkout v4.6.0
```

NEW SUPPORTED PLATFORM Nine new ELF e_machine architecture constants in the elf module

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

October 27, 2026 - Custom licenses, exclusion comments, and CVE severity (2026-10-27)

BLOG ↗ RANK

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

> Finite State shipped scan email notifications with per-scope subscriptions, programmatic custom license lifecycle management, and exclusion comment visibility on components, all with new public API endpoints.

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

01 Scan completion and failure email notifications NEW

91

Adds a `POST /api/public/v0/notifications/subscriptions` endpoint to subscribe to email notifications for every scan in a project programmatically. The user who started a scan is automatically emailed on completion or failure, and users can subscribe at three scopes — a single scan, every scan in a Folder, or every scan in a Project — including a Notify me option on a running scan to subscribe mid-flight. A new Notification Settings view lets users choose which events to receive and filter subscriptions by scope, while organization administrators can enable or disable the scan email channel org-wide under Integrations with a single switch that reflects delivery status even after a recipient has used an unsubscribe link.

Enable the scan email channel for your organization so members receive scan completion and failure emails — do this once before individual subscriptions take effect.

> 📍 In the Finite State platform, go to Integrations, find the scan email channel, and toggle the switch to enable scan emails for the organization.

— Names endpoint, scopes, UI views and admin toggle with clear setup steps

October 27, 2026 - Scan email notificat…

02 Custom license lifecycle management NEW

87

Adds `PATCH /v0/licenses` and `DELETE /v0/licenses` (with `id` as a query parameter) to edit, deprecate, restore, or delete a custom license programmatically. The same actions are available in the Settings > Policy > Licenses UI, with audit trail entries recorded for create, update, and delete.

Retire a custom license that is no longer valid without breaking policies that still reference it — the delete is refused until the policy is updated.

```
$ curl -X DELETE 'https://platform.finitestate.io/v0/licenses?id=
<license-id>' \
  -H 'Authorization: Bearer <token>'
```

— Named endpoints and runnable delete command with refusal behavior noted

October 27, 2026 - Custom licenses, exc…

03 **Exclusion comment visibility on components** NEW 86

Extends `GET /api/public/v0/components` and the component detail endpoint to return the exclusion comment, the time of exclusion, and the author who excluded the component. Each exclude or include action — including the comment and full component count — is recorded in `GET /v0/audit`, and the comment, excluding user, and timestamp are surfaced on the Component view and the Excluded list in the UI.

Audit why a component was excluded from a scan result, including who excluded it and when, before deciding to re-include it.

```
$ curl
'https://platform.finitestate.io/api/public/v0/components/<compon
ent-id>' \
  -H 'Authorization: Bearer <token>'
```

Review the full audit trail of all component exclusions and inclusions, with comments and component counts, for compliance reporting.

```
$ curl 'https://platform.finitestate.io/v0/audit' \
  -H 'Authorization: Bearer <token>'
```

— Two named endpoints with runnable calls and audit trail detail

October 27, 2026 - Custom licenses, exc…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Email Security

Material Security

Feedback ⓘ announced · Oct 28, 2026

BLOG ↗ RANK

Material Security is a SaaS platform that protects enterprise email and cloud file collaboration by detecting and remediating data exfiltration, phishing, and account takeover threats.

> Material Security shipped a GA auto-classification engine for OAuth apps with new filtering and reporting, added blast-radius and sourced-investigation visibility into risky app connections, and rebuilt its Email Identity Protection and Email Redaction reports with new outcome taxonomies, PDF export, and merged metrics.

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

01 **Rebuilt Email Identity Protection report** IMPROVED 88

Rebuilds the Email Identity Protection report to place every message into exactly one of three mutually exclusive outcomes — delivered without verification, held before delivery, or blocked — with a secondary breakdown of held-message outcomes (retrieved, reported unrequested, failed verification, never attempted, verification outstanding). Separates activity by message type (password resets, verification codes, account activations, account access messages), includes a service matrix showing each service's setting (Protect , Block , Detect Only , or Disabled), and adds PDF export for leadership reviews.

Produce a leadership-ready PDF summary of Email Identity Protection coverage and held-message outcomes.

📍 In the Material console, open the Email Identity Protection report, verify the service matrix reflects the intended per-service setting (Protect / Block / Detect Only / Disabled) for each mail type, then click Export to PDF.

– Full outcome taxonomy, matrix and export path named with clear mechanism

launch-20261028-e2bbc953

02 **Auto-classification for OAuth apps** NEW 86

Adds the Auto-classify OAuth apps setting (off by default) that automatically applies the OAuth Agent's recommended classification to new and audited app connections and triggers configured remediations without manual review. Adds a Classification State filter and column (values: Suggested or Applied) to the OAuth apps list and CSV exports, plus a new OAuth Agent activity report summarizing how apps have been classified, which have been remediated, and which require further investigation.

Enable automated OAuth classification so the agent immediately remediates newly-connected malicious apps without analyst queuing.

> 📍 In the Material console, go to OAuth Agent settings and enable `Auto-classify OAuth apps`. After saving, Material surfaces a count of apps still holding `Suggested` classifications — click through and resolve them. Going forward, filter the OAuth apps list by Classification State = Suggested to surface any app awaiting a human decision.

– Names exact config key, filter values and report; remediation scope detail is thin

`launch-20261028-e2bbc953`

---

03 **Blast radius visibility for OAuth apps** NEW 70

Adds a `Blast radius` card on each OAuth app's detail view showing how much sensitive content (files and messages) the app's connected accounts can reach, with links from either count into the Explorer view to inspect the specific at-risk content.

Quantify the exact data exposure a risky OAuth app has before deciding whether to revoke its tokens.

> 📍 Open the OAuth app's detail view and review the `Blast radius` card, which shows the count of sensitive files and messages reachable by the app's connected accounts. Click either count to open the Explorer view and inspect the specific content at risk.

– Names card and navigation path but no scoring methodology or limits

`launch-20261028-e2bbc953`

---

04 **Vendor trust and scope risk columns** NEW 60

Adds `Vendor Trust Level` and `Scope Risk Level` columns to the OAuth Security Issues view for sorting and filtering apps by reputation and permission breadth.

– Names exact columns and view but no scoring methodology explained

`launch-20261028-e2bbc953`

---

05 **Unified Email Redaction report** IMPROVED 60

Replaces the previous Email Redaction and Sensitive Content reports with a unified Email Redaction report that includes an estimate of currently redacted messages with a weekly trend, histograms of access requests and retrieved-message age, and a chart of access requests over time.

– Names merged reports and metrics but no navigation path given `launch-20261028-e2bbc953`

THINNER COVERAGE BELOW

06 **Sourced OAuth app investigation verdicts** IMPROVED 40

OAuth app investigations now show the web searches and sources the OAuth Agent used to reach its verdict, with every claim linking back to its source.

— Describes behavior but no UI location or config surface named `launch-20261028-e2bbc953`

WAS THIS USEFUL? 👍 👎

Sublime Security

Feedback ⓘ

Platform v1.90.133 (Attack Score Flash) (2026-10-27) BLOG ↗ RANK

Sublime Security is an email security platform that uses detection-as-code rules to identify and block phishing, business email compromise, and other malicious messages.

> Sublime Security made Attack Score Flash generally available, bringing ML-based threat verdicts to all inbound mail and wiring it into existing automations and the Autonomous Security Analyst.

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

01 **Attack Score Flash general availability** NEW 71

Attack Score Flash now runs on all inbound mail, producing high-precision Malicious, Suspicious, Spam, and Graymail verdicts to catch messages existing rules may miss, similar to Attack Score. Potentially Malicious and Spam detections are promoted into the core feed so custom rules are no longer needed to surface them. Existing automations that trigger on Attack Score verdicts now also act on messages flagged by Attack Score Flash, increasing automated flagged-message volume, and the Autonomous Security Analyst automatically reviews Attack Score Flash Suspicious verdicts via the existing Send suspicious messages to ASA automation. Potentially Malicious detection is Sublime Enterprise-only, while Attack Score Flash itself is available on both Sublime-managed cloud and self-hosted deployments.

— Clear mechanism and scope but no UI path or config key given

`Platform v1.90.133 (Attack Score Flash)…`

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Aviatrix

Feedback ⓘ announced · Oct 28, 2026 PRESS ↗ RANK

Welcome to the comprehensive documentation for Aviatrix's cloud networking platform. Here you'll find everything you need to deploy, configure, and manage your multi-cloud networking infrastructure.

> Aviatrix launched Clearance, a per-workload outbound and lateral containment product within its Containment Platform, paired with an autonomous Active Defense policy agent, Shadow Rule Set monitor-mode testing, one-action enforcement/backout, and a free assessment plus limited-time free enterprise access via AWS Marketplace.

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

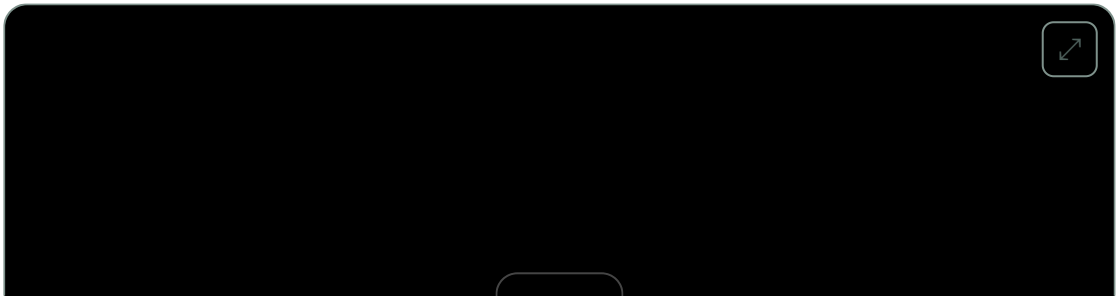
01 **Containment Assessment and pricing tiers** NEW 83

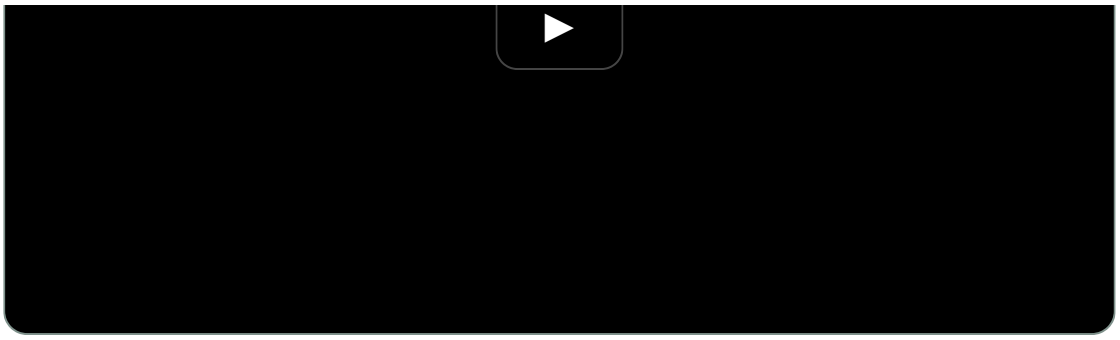
A free read-only Containment Assessment at `aviatrix.ai/containment-assessment` installs nothing and can reserve AWS Marketplace enrollment. Clearance is free for all enterprises — covering Outbound Containment, Lateral Containment, and Intrusion Prevention with no workload cap — from November 1, 2026 through May 1, 2027 via AWS Marketplace. A General Availability tier offers the Containment Assessment plus 15 workloads in monitor or enforce mode with no time limit, priced per workload and delivered on Aviatrix Cloud (same controls also available on Aviatrix Enterprise). Aviatrix Active Defense is available today in free early access for design partners.

— Exact URL, dates, workload limits and tiers named launch-20261028-5ed5a331

02 **Outbound and Lateral Containment in Clearance** NEW 74

Aviatrix Clearance™, part of the Aviatrix Containment Platform, introduces Outbound Containment (limits each workload to outbound destinations it is explicitly cleared for, named by domain rather than network address) and Lateral Containment (restricts workload-to-workload reachability by identity across clouds and Kubernetes/serverless paths not governed by hub-based inspection). It enforces per-workload outbound and lateral communication policy by workload identity and protocol across clouds, down to individual containers, serverless functions, databases, and AI agents.





— Names mechanism and scope but no exact config or command launch-20261028-5ed5a331

THINNER COVERAGE BELOW

03 **Active Defense autonomous policy agent** NEW 53

Aviatrix Active Defense™ is an autonomous agent layer that models the estate from live traffic, drafts per-application policies with no tagging project or inventory import, proves each policy in monitor mode on real traffic, and keeps policies current as the estate changes.

— Explains behavior but no interface or command given launch-20261028-5ed5a331

04 **Shadow Rule Set monitor mode** NEW 47

Shadow Rule Set monitor mode evaluates a new policy alongside the live policy on real traffic and shows exactly which connections would be blocked before any enforcement change is made.

— Clear mechanism but no named flag or command launch-20261028-5ed5a331

05 **One-action enforcement and backout** NEW 27

Delivers single-setting enforcement and single-action backout for every containment policy.

— Bare description, no UI path or setting name given launch-20261028-5ed5a331

WAS THIS USEFUL?

◆ IDS / IPS

CrowdSec

Feedback ⓘ v1.8.2 NOTES ↗ RANK

Open-source IDS/IPS, WAF and bot detection for Linux, Windows, Docker and Kubernetes, with a crowdsourced blocklist of malicious IPs.

> CrowdSec v1.8.2 adds WAF rate limiting and a SkipProcessing() bypass helper, expands bot detection with custom JS fingerprinting and a new cscli bot subcommand, and ships several new expression helpers plus observability and LAPI improvements.

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

01 **New expression helpers** NEW 88

Adds `GetDecisionsSinceCountByScenario` expression helper for querying per-scenario decision counts, usable in postoverflow or rule expressions to gate actions on how many decisions a scenario has already produced in a time window, e.g. `GetDecisionsSinceCountByScenario('crowdsecurity/http-probing', '1h') > 5`, and `LookupHost` DNS cache support in expression helpers to reduce repeated DNS lookups.

Use `GetDecisionsSinceCountByScenario` in a postoverflow or rule expression to gate actions on how many decisions a scenario has already produced in a time window.

```
$ GetDecisionsSinceCountByScenario('crowdsecurity/http-probing', '1h') > 5
```

— Named helper with a concrete runnable expression example. Named helper but no usage example or config detail given.

v1.8.2

02 **SkipProcessing() helper for WAF rule bypass** NEW 85

Adds SkipProcessing() to the WAF for conditionally bypassing rule evaluation, letting a matching request short-circuit all further rule evaluation — useful for allowlisting trusted internal scanners.

Call SkipProcessing() inside a WAF rule to short-circuit all further rule evaluation for a matching request, useful for allowlisting trusted internal scanners.

```
$ SkipProcessing()
```

— Exact function name with a runnable example showing its use. v1.8.2

03 **Bot detection CLI, custom fingerprinting and cookie flags**

NEW

78

Adds a `cscli bot` subcommand for managing bot detection from the CLI, a FlagSubmission() helper for marking submission events, custom JS fingerprinting allowing user-defined and hub-distributed detection rules, and optional `SameSite=None` and `Partitioned` flag support on bot detection cookies.

Inspect or manage bot detection state directly from the CLI using the new bot subcommand.

```
$ cscli bot
```

– Names exact CLI command and cookie flags; fingerprinting mechanism only briefly described.

v1.8.2

04 **New Prometheus metrics for features, alerts and CAPI sync**

NEW

63

Adds a `cs_features` metric to track which Security Engine features are active, a `kind` label on the `cs_alerts` metric, and a CAPI last-pull timestamp plus error counter to metrics.

– Names exact metric and label names; no query examples given. v1.8.2

THINNER COVERAGE BELOW

05 **LAPI decision stream concurrency limit and paging** IMPROVED

50

Adds an optional concurrency limit for simultaneous decision stream requests via LAPI and makes the decision stream page size configurable.

– Options described generically without exact config key names. v1.8.2

06 **Category field on AppSec request score** NEW

43

Adds an optional `category` field to the AppSec request score for finer-grained scoring classification.

– Field named but no example of configuring or reading it. v1.8.2

07 **Automatic TLS certificate reload** IMPROVED

42

Enables automatic TLS server certificate reload when certificate files change on disk.

– Behavior described but no config flag or path named. v1.8.2

08 **Request-rate-based blocking in the WAF** NEW

40

Adds a rate limiting feature to the WAF enabling request-rate-based blocking at the application layer.

— No config key or flag named for enabling the limit. v1.8.2

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Google distroless

Feedback ⓘ changes since 2026-10-26

CODE ↗

Language focused docker images, minus the operating system.

UPDATED SUPPORTED PLATFORM Bumping non-snapshot packages to latest stable versions

└──▷ BREAKING ON UPGRADE

! The `--set-github-output-on-diff` flag has been removed from `diff.bash`; GitHub Actions output is now written automatically when `GITHUB_ACTIONS=true` and `GITHUB_OUTPUT` is set.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ

Catch leaked AWS keys before someone runs up your Bedrock bill: NOTES ↗ RANK

Plerion is a cloud security platform that detects misconfigurations, vulnerabilities, and threats across cloud environments and provides remediation guidance.

> Plerion's latest release adds LLMjacking and compromised-key detections for AWS, a new Azure Function App end-of-support check, and widens access-logging coverage across API Gateway, load balancers, VPCs and EKS, while retiring one EFS logging check.

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

01 LLMjacking and compromised AWS key detections NEW

74

Adds `PLERION-CLOUDTRAIL-84` (Critical) to detect compromised AWS keys via AWS Health exposed/compromised access key reports, naming the key and IAM user (requires `us-east-1` in your CDR regions); `PLERION-CLOUDTRAIL-85` (High) to detect deletion of Amazon Bedrock model invocation logging, a LLMjacking cover-track technique; and `PLERION-AWS-1252` (Low) to check whether CloudTrail logs data events for `AWS::BedrockMantle::Project`, the endpoint for OpenAI-compatible and Anthropic Messages APIs, via an advanced event selector on a multi-region trail.

— Names three checks and exact mechanism, but detection is passive, not user-driven.

Catch leaked AWS keys before someone ru…

02 Azure Function App end-of-support runtime check NEW

73

Adds `PLERION-AZURE-463` (Medium) to detect Azure Function App runtimes past Microsoft's end-of-support date, covering .NET, .NET isolated, Java, Node.js, PowerShell and Python; flags Python 3.8/3.9, .NET 6, PowerShell 7.2 and Node.js 20 on next scan.

— Names exact runtimes flagged, giving a clear remediation target.

Catch leaked AWS keys before someone ru…

03 Access logging filter expanded to more AWS asset types

IMPROVED

70

Expands the Access logging filter on the Assets page and the matching custom-report field to cover API Gateway (`PLERION-AWS-494` , `PLERION-AWS-799` , `PLERION-AWS-801`), load balancers (`PLERION-AWS-791` , `PLERION-AWS-792` , `PLERION-AWS-795`), VPCs (`PLERION-AWS-37`) and EKS clusters (`PLERION-AWS-107`).

— Names every check ID and the exact UI filter to use them from.

Catch leaked AWS keys before someone ru…

04 **Widened pass criteria for CloudFront and Key Vault logging checks**

IMPROVED 68

Widens `PLERION-AWS-231` CloudFront general-logging check to pass distributions whose standard logging (v2) delivers via a CloudWatch Logs delivery to CloudWatch Logs, Amazon S3 or Amazon Data Firehose, and widens `PLERION-AZURE-74` Key Vault logging check to pass vaults whose diagnostic setting enables the `audit` or `allLogs` category group, not only an explicit `AuditEvent` selection.

— Explains exact before/after logic for two named checks.

Catch leaked AWS keys before someone ru…

THINNER COVERAGE BELOW

05 **GCP Cloud Function credential scanning promoted to active finding**

IMPROVED 57

Promotes `PLERION-GCP-18` (Medium) from a manual recommendation to an active finding: scans Cloud Function runtime and build environment variables for credentials (API keys, access tokens, private keys), naming the variable but never its value.

— Clear behaviour change but no user-facing action beyond review.

Catch leaked AWS keys before someone ru…

06 **EFS access logging check retired** BREAKING 45

`PLERION-AWS-835` (EFS access logging check) is retired; its passed findings are removed on the next scan.

— Names the retired check and effect but no migration guidance.

Catch leaked AWS keys before someone ru…

07 **Pleri AI assistant names vulnerable file or package path**

IMPROVED 35

Pleri (AI assistant) now names the specific file or package path carrying a CVE on each affected asset.

— Describes the change but no mechanism or scope beyond one sentence.

Catch leaked AWS keys before someone ru…

└──▷ **BREAKING ON UPGRADE**

! `PLERION-AWS-835` (EFS access logging) is retired; its passed findings are removed on the next scan.

WAS THIS USEFUL? 👍 👎

◆ **Runtime Security**

aya

Feedback ⓘ `aya-ebpf-macros-v0.4.0` NOTES ↗ RANK

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

> Aya's v0.4.0 macros release adds BTF-based map pinning and validates pinned maps before reuse, tightening how maps are shared across the BPF filesystem.

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

01 **BTF map pinning by name with reuse validation** NEW

60

`aya-ebpf-macros` v0.4.0 adds support for BTF map pinning by name in `aya-ebpf`, letting maps be pinned to the BPF filesystem using their BTF-derived names, and `aya` now validates pinned maps before reuse by checking type and size correctness when attaching to an existing pinned map.

— Names the mechanism and crates but no exact API or flag `aya-ebpf-macros-v0.4.0`

WAS THIS USEFUL? 👍 👎

Alibaba Cloud OpenSandbox

Feedback ⓘ release-1.2.0

NOTES ↗ RANK

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

> OpenSandbox 1.2.0 adds auto-pause/wake-on-access lifecycle policies, a Credential Vault API, OCI registry snapshot publishing, and a diagnostics context endpoint across all five SDKs, alongside a structured renewal-policy field that replaces a deprecated legacy key and several hardening fixes for gVisor, podman, and Kubernetes egress.

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

01 **Auto-pause and idle sweeper controls for sandboxes** NEW

95

Adds `lifecycle.idlePolicy` create field with `idleTimeoutSeconds` (30–86400), `onIdle: pause`, `wakeOnAccess`, and optional `activeProbe` to enable auto-pause and wake-on-access for fast sandboxes; server-side config moves under `[automation.idle_pause]`, requiring `redis.dsn` when the sweeper is enabled, and the idle sweeper now also supports an `OpenSandbox-Skip-Activity` sentinel header to suppress one activity observation, stripped on both gateway and proxy hops. All five SDKs gain typed client support with fail-fast validation and a wake-retry contract: on a `503` with `OpenSandbox-Wake: expired`, SDKs retry after `Retry-After` up to a bounded budget (default 5 attempts / 60s) and raise a typed wake-timeout exception on exhaustion, with stream-start wake-retry active in JavaScript, Python, and Go.

Create a fast sandbox that auto-pauses after 5 minutes of idle and wakes transparently on the next access, keeping costs low for infrequently used environments.

```
$ curl -X POST https://api.opensandbox.example/v1/sandboxes \
  -H 'Content-Type: application/json' \
  -d '{
    "template": "my-template",
    "lifecycle": {
      "idlePolicy": {
        "idleTimeoutSeconds": 300,
        "onIdle": "pause",
        "wakeOnAccess": true
      }
    }
  }'
```

Enable the MITM CA pinning feature so that egress sidecar restarts never rotate the trust anchor under live sandboxes in a Kubernetes deployment.

```toml
[kubernetes]
template_mitm_ca_secret = "mitm-ca-secret"

[automation.idle_pause]
enabled = true

[redis]
dsn = "redis://localhost:6379/0"
```

— Full mechanism, config keys, retry contract, and runnable example given. Names exact header and where it is stripped.

release-1.2.0

02 **MITM CA pinning for Kubernetes sandbox templates** NEW

86

Adds `[kubernetes] template_mitm_ca_secret` server config (default off) to bake a pinned MITM CA into server-created SandboxTemplates via `spec.egressMitmCertificate`, so an egress sidecar restart no longer rotates the trust anchor under live sandboxes.

Enable the MITM CA pinning feature so that egress sidecar restarts never rotate the trust anchor under live sandboxes in a Kubernetes deployment.

```toml
[kubernetes]
template_mitm_ca_secret = "mitm-ca-secret"

[automation.idle_pause]
enabled = true

[redis]
dsn = "redis://localhost:6379/0"
```

— Exact config key with a runnable config example. release-1.2.0

03 **Structured renewal policy replaces legacy extension key**

86

BREAKING

Adds `lifecycle.renewalPolicy.extendSeconds` (300–86400) structured field (OSEP-0009), replacing the opaque `extensions['access.renew.extend.seconds']` key; the legacy key is deprecated

through 1.x and now rejected with `400` instead of silently accepted as a no-op, and will be removed in 2.0.

— Names old and new fields plus a concrete migration deadline. release-1.2.0

04 **Credential Vault API for fast sandboxes** NEW 83

Adds Credential Vault support via a `POST|GET|PATCH|DELETE /v1/sandboxes/{id}/credential-vault` API, with credential values sanitized on every public read and `expectedRevision` used as the cross-replica PATCH backstop; a Credential Vault client contract is added to all five SDKs that routes vault calls through the server for template-backed sandboxes.

— Exact endpoint and methods named, but no runnable example. release-1.2.0

05 **Diagnostics context endpoint** NEW 78

Adds `GET /v1/sandboxes/{sandboxId}/diagnostics/context` endpoint returning sandbox ID, lifecycle state, backend type, and backend-defined identifiers for Docker, Kubernetes, and Fast Sandbox backends without connecting to execd; unsupported extension backends return `501`.

— Exact endpoint, response fields, and error code given. release-1.2.0

06 **gVisor snapshot safety preflight** NEW 73

Adds preflight rejection of gVisor snapshots unless effective args include `--overlay2=none`, preventing silent write-drop on `docker commit` under gVisor's default `--overlay2=root:self`.

— Exact flags named; concrete failure mode explained. release-1.2.0

07 **Podman network mode passthrough** NEW 60

Passes through podman `pasta`, `slirp4netns`, and `none` network modes, with `networkPolicy` rejected since there is no bridge to attach to.

— Names exact modes and rejection condition. release-1.2.0

THINNER COVERAGE BELOW

08 **Unbounded SDK file transfers** IMPROVED 55

Removes whole-call timeouts from SDK file transfers in all five languages, eliminating the 30s default cap that previously caused large or slow transfers to fail mid-flight.

— Explains before/after but no API or flag named. release-1.2.0

09 **OCI registry snapshot publishing** NEW 45

Supports publishing Fast Sandbox snapshots and template builds to, and booting from, OCI registries via OverlayBD OCI lifecycle.

— Describes capability but no concrete command or config key. release-1.2.0

└──▷ **BREAKING ON UPGRADE**

! The `extensions['access.renew.extend.seconds']` key is deprecated and fast sandboxes now reject it with 400 instead of silently accepting a no-op; migrate to `lifecycle.renewalPolicy.extendSeconds` before upgrading to 2.0, where the legacy key is removed.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ **Endpoint Security (EDR/XDR)**

Huntress

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗

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

Huntress adds Identity Posture Assessment (on-demand PDF + API) for M365 tenants and Lookback for 180-day historical login analysis on ITDR onboarding.

Adds Identity Posture Assessment available via the Huntress API, enabling automated retrieval and integration of M365 tenant posture reports into existing workflows.

Adds Identity Posture Assessment, an on-demand PDF report evaluating Microsoft 365 tenants against the Huntress Identity Security Framework — showing gaps, misconfigurations, severity-ranked next steps, and co-brandable output for client QBRs and onboarding.

NEW DETECTION Adds Lookback, which automatically runs Huntress Unwanted Access detections across up to 180 days of Microsoft 365 login history when an eligible tenant onboards to Managed ITDR, surfacing earlier compromises and active threats from day one.

Lookback findings appear as a dedicated section of the Identity Security Assessment, delivered by email and available in the Huntress portal within 48 hours of onboarding.

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Agentless Vulnerability & Threat Detection October 2026 update (2026-10-26)

BLOG ↗ RANK

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

> Trend Vision One's Agentless Vulnerability & Threat Detection can now be deployed directly inside a customer's own AWS VPC via CloudFormation, giving teams tighter network isolation and control over scanning traffic.

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

01 **In-VPC deployment for Agentless Vulnerability & Threat Detection**

NEW 65

Agentless Vulnerability & Threat Detection can now be deployed directly into a customer's AWS Virtual Private Cloud (VPC) via CloudFormation, enabling greater network security control and resource isolation. To use it, go to Cloud Security → Cloud Accounts → AWS, select the AWS account, and choose the CloudFormation deployment method to deploy into the VPC.

Deploy Agentless Vulnerability & Threat Detection into your own VPC to enforce network isolation and control over scanning traffic.

> 📍 In Trend Vision One, go to Cloud Security → Cloud Accounts → AWS, select your AWS account, and choose the CloudFormation deployment method to deploy Agentless Vulnerability & Threat Detection into your VPC.

— Names deployment method and exact UI navigation, but no mechanism detail

Agentless Vulnerability & Threat Detect…

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗ RANK

Tines lets teams put AI to work across their most sensitive systems, data, and processes, with the security and governance to do it safely.

> Tines 44.0.0 removes the legacy records migration system and its associated configuration keys ahead of a hard cutover, while adding keyless cloud credentials, fallback/custom AI model settings, a visual user-group mapping builder, and pagination and sorting upgrades to Record Lists.

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

01 Removal of legacy records migration system

BREAKING

95

Tines 44.0.0 removes the `RECORD_RESULT_SET_TABLE_NAME` and `RECORDS_MIGRATION_DISABLED` configuration keys, which must be deleted from your configuration before upgrading — they no longer have any effect on records. The release also strips out all legacy records migration code including repair tools, status banners, upgrade checks, dual writing, and partition maintenance; upgrading with unmigrated records makes those records invisible, and the old table is retained only until 45.0.0, after which it will be dropped. Do not upgrade if admin banners show 'Data migration in progress,' 'Verifying data integrity...,' discrepancies, or `Some records are unavailable` — instead run `tines:verify_records_migration` from a worker and wait for `SAFE TO UPGRADE` before proceeding.

Verify records migration is complete before upgrading from 41.x to 44.0.0 — run this from a worker and wait for `SAFE TO UPGRADE` before proceeding.

```
$ tines:verify_records_migration
```

— Names exact config keys and verification command with before/after behaviour.

changelog-20261028-1ea45807

02 **Fallback models and custom settings for AI actions** NEW

65

Adds fallback models for AI Agent actions so a secondary model is used automatically when the primary is unavailable, adds custom model settings for AI providers, and adds `Claude Haiku 5.5` as an available model for AI actions.

– Names fallback mechanism and a specific model, but no config path.

changelog-20261028-1ea45807

THINNER COVERAGE BELOW

03 **Record Lists: new field types, pagination and sorting**

IMPROVED

55

Adds new field types for records, and adds cursor pagination, multi-field sorting, and lighter queries to Record List actions.

– Names several concrete capabilities but no field types or limits.

changelog-20261028-1ea45807

04 **Keyless cloud credentials via workload identity federation**

NEW

40

Adds keyless cloud credentials via workload identity federation.

– Names the mechanism but gives no setup steps or config surface.

changelog-20261028-1ea45807

05 **Story change control covers action attributes** IMPROVED

35

Extends story change control requirements to cover action attributes.

– Short statement without specifics on which attributes are covered.

changelog-20261028-1ea45807

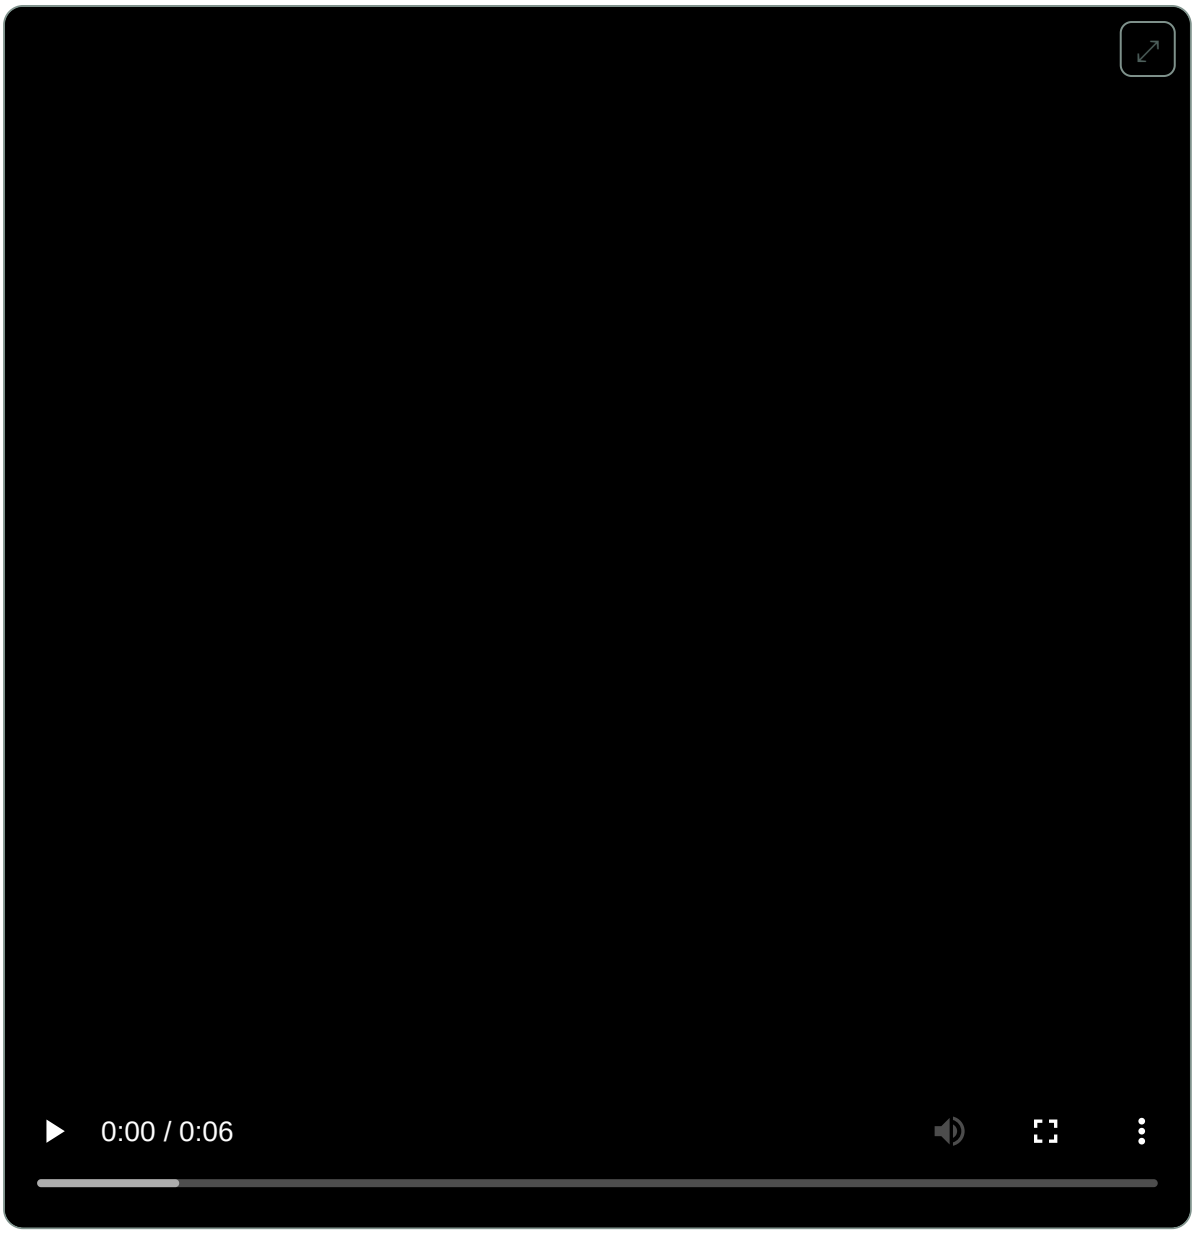
06 **Visual builder for user group mappings** NEW

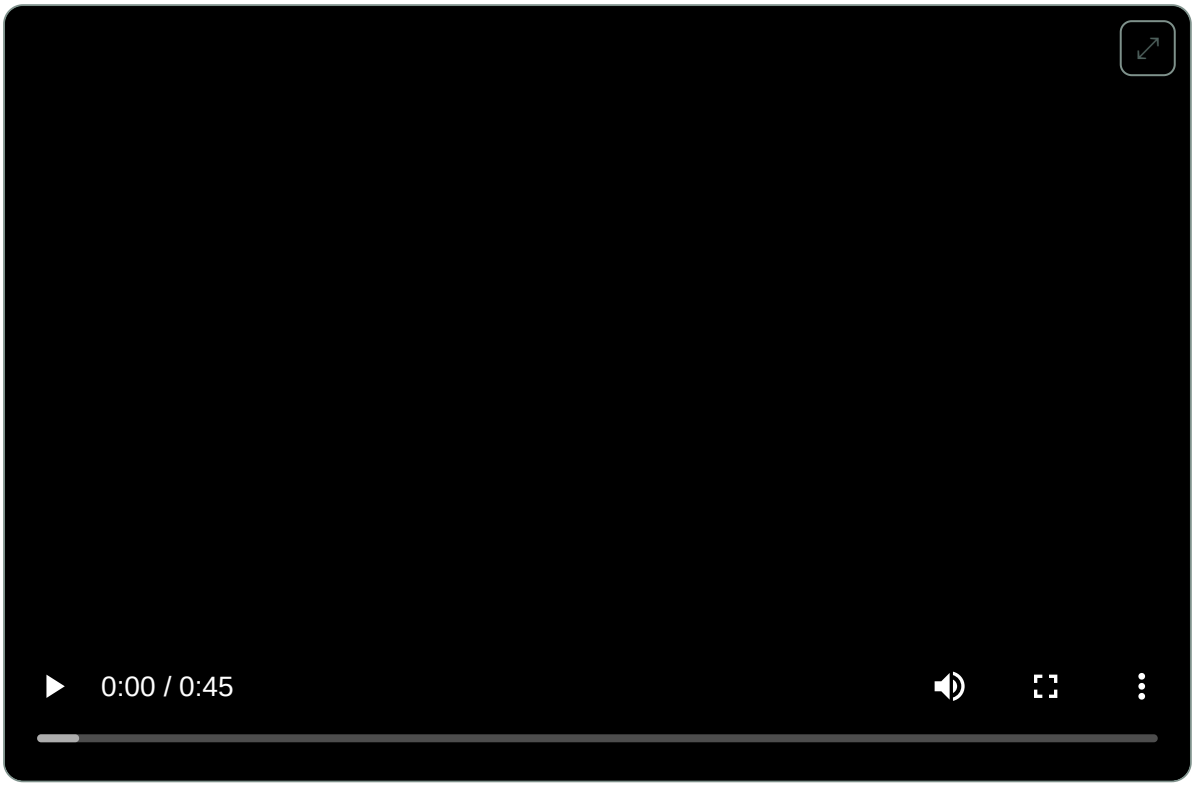
30

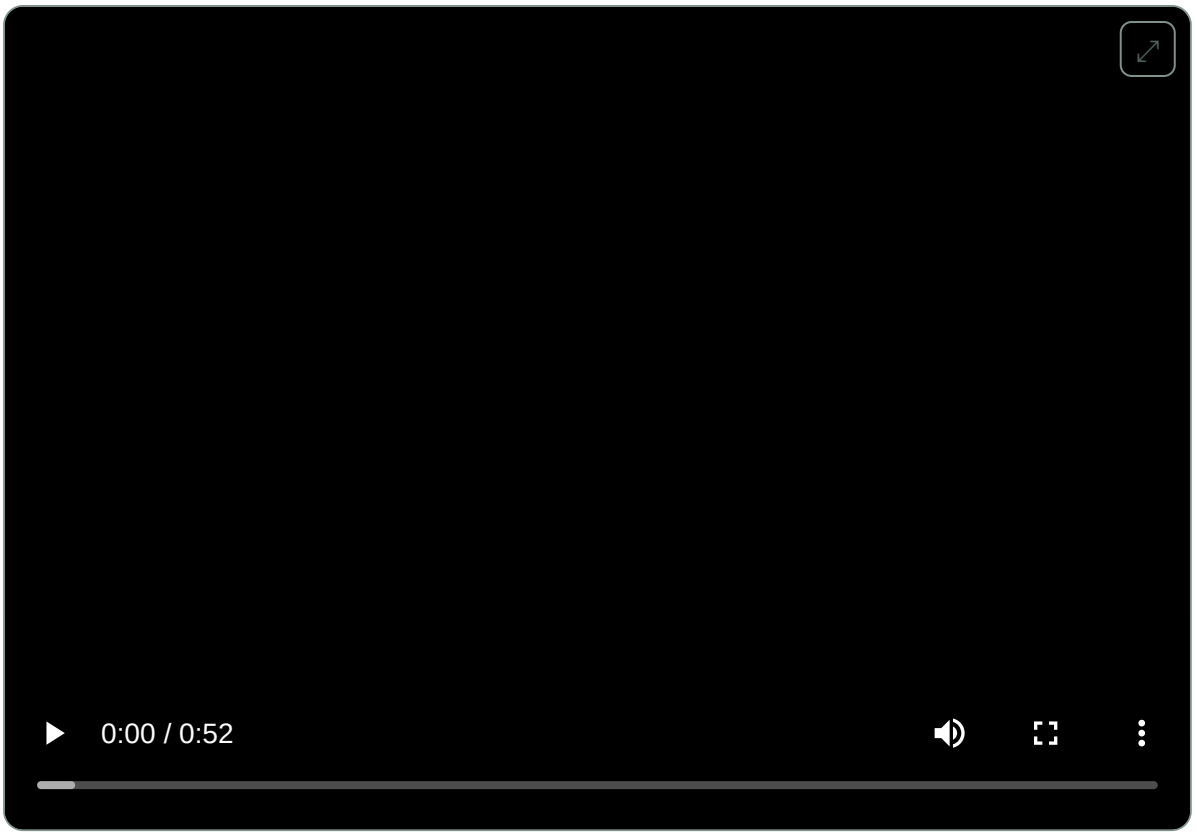
Adds a visual builder to configure user group mappings.

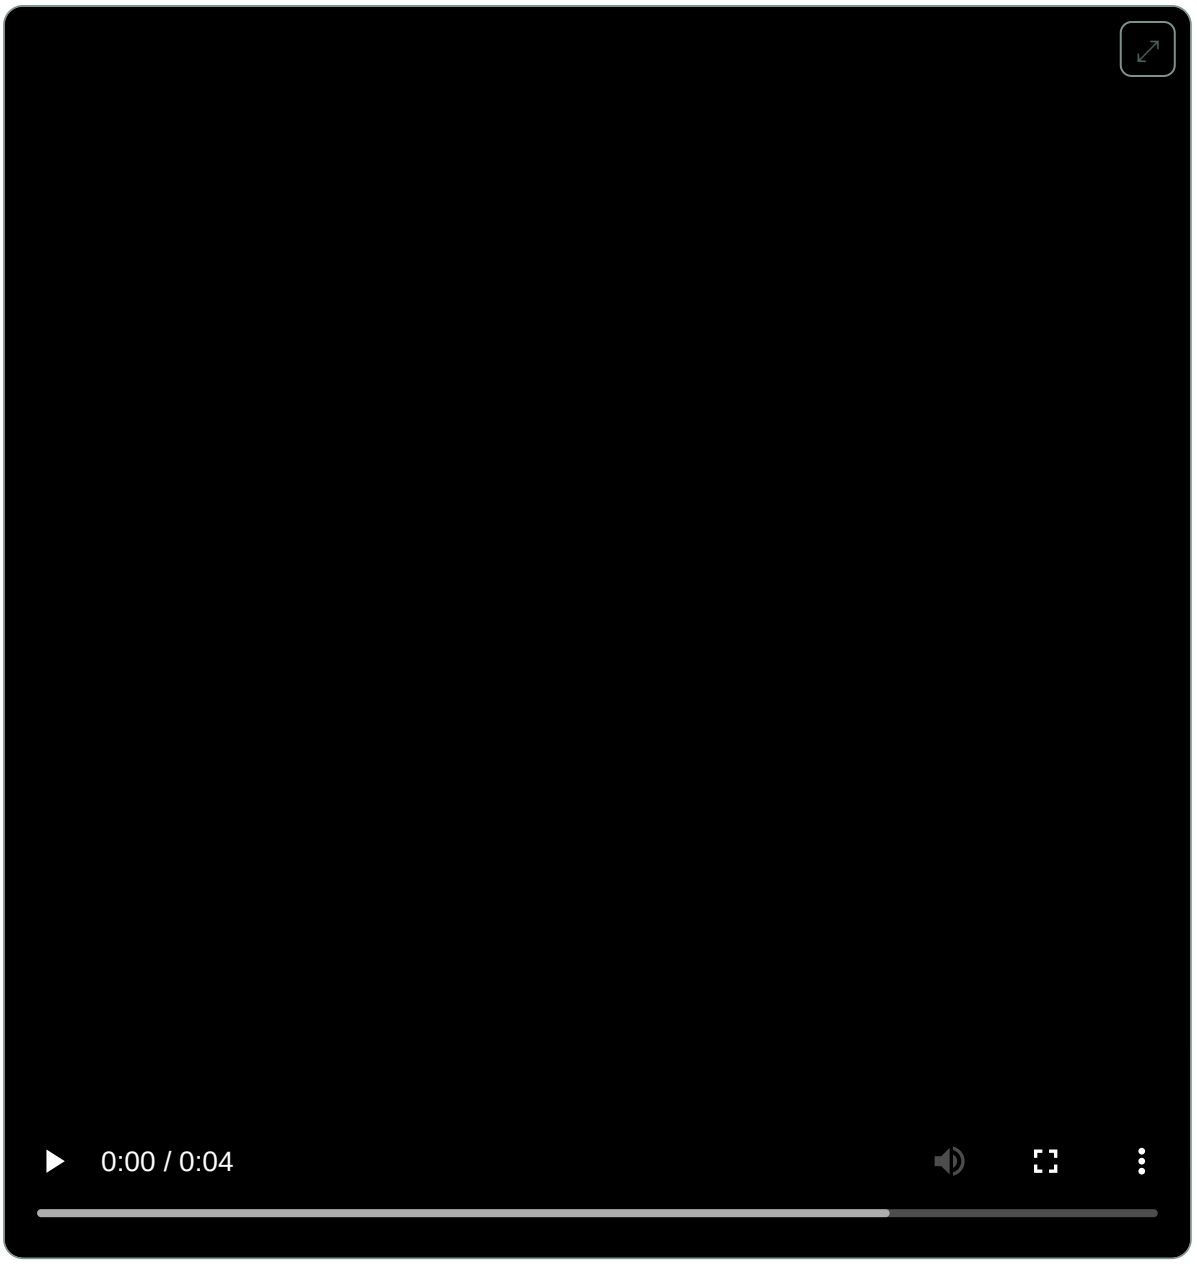
– Bare description with no mechanism or UI location given. changelog-20261028-1ea45807

└──▷ ALSO FROM THESE RELEASES


0:00 / 0:06


0:00 / 0:45


0:00 / 0:52



└──▷ **BREAKING ON UPGRADE**

! `RECORD_RESULT_SET_TABLE_NAME` and `RECORDS_MIGRATION_DISABLED` are removed and must be deleted from your configuration — they no longer have any effect on records.

! 44.0.0 removes all legacy records migration code including repair tools, status banners, upgrade checks, dual writing, and partition maintenance. Upgrading with unmigrated records makes those records invisible; the old table is retained only until 45.0.0, after which it will be dropped. Do not upgrade if admin banners show 'Data migration in progress,' 'Verifying data integrity...,' discrepancies, or 'Some records are unavailable.'

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Beelzebub

Feedback ⓘ v3.10.0 NOTES ↗

A secure low code deception runtime framework, leveraging AI for System Virtualization.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.10.0 https://github.com/beelzebub-labs/beelzebub.git
# already have the repo? check out this version:
$ git checkout v3.10.0
```

NEW DETECTION WordPress 6.9.4 service lure

NEW DETECTION Ollama service lure

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

AI Functions Unity Catalog permissions is now generally available (2026-10-27)

BLOG ↗ RANK

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

> Databricks made AI Functions permissioning and account-level security features generally available this window, both now active by default without manual opt-in.

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

01 **Account-level security features on by default at GA** IMPROVED 60

Account-level security features covering network connectivity (private connectivity), network security (ingress/egress controls), encryption, and enhanced security and compliance are now automatically available to all accounts upon General Availability, removing the requirement for account admins to manually enable them on the Previews page.

– Names the feature areas and removed manual step, no further config detail

Account-level security features are now…

THINNER COVERAGE BELOW

02 **Unity Catalog permissions for AI Functions** NEW 55

Metastore admins can now apply Unity Catalog permissions on the `system.ai` schema to control which task-specific AI Functions users are allowed to run. This capability is now generally available.

– Names the schema and control point but no further mechanism

AI Functions Unity Catalog permissions …

NEW INTEGRATION Default OAuth client IDs for external AI agents connecting to Databricks MCP servers

> Default OAuth apps for external AI agents are now generally available (2026-10-27)
>
> BLOG ↗ LATEST STABLE
>
> External AI agents can now authenticate to Databricks MCP servers via default OAuth client IDs, eliminating manual OAuth app setup by account admins.

> Enables ChatGPT, Codex, Claude, Claude Code, and Cursor to connect to Databricks MCP servers (including the Genie One MCP server) using default OAuth client IDs, with no OAuth app creation required from an account admin.

WAS THIS USEFUL?  

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Security Validation (BAS & Automated Pentest)

Horizon3 NodeZero

Feedback ⓘ announced · Oct 28, 2026

PRESS ↗ RANK

Horizon3 nodezero is a horizon3.ai enables organizations to continuously assess the security posture of their enterprise across many attack surfaces.

> NodeZero shipped a certified Cortex XSOAR integration that pushes high- and critical-severity findings into XSOAR as actionable incidents, paired with a new `1-Click Verify` action to retest remediated findings without leaving the XSOAR workflow.

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

01 Cortex XSOAR integration for NodeZero findings NEW

75

A certified NodeZero integration is now available on the Cortex Marketplace that automatically ingests NodeZero high- and critical-severity open findings into Cortex XSOAR as actionable incidents, complete with affected asset data, exploitation evidence, and remediation guidance. It supports continuous Find-Fix-Verify-Repeat workflows: remediated findings stop returning as open incidents in XSOAR, while unresolved exploitable weaknesses continue to surface for action. Install via Cortex XSOAR Marketplace > search NodeZero > Install, then configure NodeZero API credentials under Settings > Integrations > NodeZero.

After installing the NodeZero integration from the Cortex Marketplace, route incoming NodeZero incidents through an existing XSOAR playbook for enrichment and remediation assignment.

📍 In Cortex XSOAR, go to Marketplace › search NodeZero › Install. Then open Settings › Integrations › NodeZero and configure your NodeZero API credentials. New high- and critical-severity NodeZero findings will appear automatically as incidents under Incidents › NodeZero Findings.

After NodeZero completes an autonomous pentest, have XSOAR playbooks automatically receive and route the exploitable findings — then retest a patched asset without leaving the XSOAR workflow.

📍 In Cortex XSOAR, go to Marketplace › search NodeZero › Install the certified NodeZero integration. Configure the integration with your NodeZero credentials, then use an existing or new playbook to consume incoming NodeZero incidents. After patching, trigger 1-Click Verify on the incident to confirm the finding is no longer exploitable.

— Names marketplace, incident fields, and exact install navigation.

launch-20261028-b7d7cdab launch-20261027-b7d7cdab

02 **`1-Click Verify` retest action in XSOAR** NEW 65

Enables `1-Click Verify` from within the remediation workflow in Cortex XSOAR to retest eligible findings and confirm whether they remain exploitable after a fix is applied, without re-running a full pentest. Triggered from an open NodeZero incident via Actions > 1-Click Verify, which updates the incident status based on the retest result.

After applying a fix, use 1-Click Verify to retest a specific finding directly from Cortex XSOAR without re-running a full pentest.

> 📍 In Cortex XSOAR, open the NodeZero incident for the remediated finding › click Actions › select 1-Click Verify to trigger NodeZero to retest only that finding and update the incident status based on whether it remains exploitable.

— Names the exact action and UI steps to trigger it. launch-20261028-b7d7cdab launch-20261027-b7d7cdab

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Censys

Feedback ⓘ announced · Oct 27, 2026 BLOG ↗ RANK

Censys is a internet-wide scanning and search platform that discovers and monitors hosts, services, and certificates to help security teams assess their exposure.

> Censys turned its Investigator threat-hunting workflow into a full API, letting scripts, SOAR playbooks, and AI agents submit indicators, upload raw intel documents, and retrieve full campaign-cluster reports programmatically.

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

01 Investigator API for bulk indicator investigations NEW

95

Censys Investigator is now exposed as a programmable API. `POST /files` uploads raw HTML/text (e.g. a vendor blog post) as investigation input and returns a `file_id` for use in subsequent investigation calls. The investigation endpoint accepts up to 100 IPs, domains, or file IDs per job with optional `start_time` and `end_time` window parameters, returning a downloadable archive containing a CenQL cluster query and full infrastructure report. `GET /usage` checks an organization's remaining weekly investigation allowance and `next_expire_time` before submitting jobs. Each organization receives 150 Investigator API investigations per week (up to 3 concurrent per user), with results retained for 7 days and no Censys Credits consumed.

Upload a vendor threat-intel report and expand its indicators into a full campaign cluster, then save the CenQL query as a watched Collection.

```
$ # 1. Upload the report HTML as a file
curl -X POST 'https://api.platform.censys.io/files?org_id=
<org_id>' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"content_type": "text/html", "filename": "vendor-
report.html", "size_bytes": 42800}'

# 2. PUT the actual bytes to the returned upload_url, then run
investigation
#    (investigate() wraps the investigation submission call with
file_ids=["<file_id>"], start_time="...")

# 3. Turn the returned CenQL into a persistent Collection
curl -X POST 'https://api.platform.censys.io/v3/collections?
org_id=<org_id>' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"name": "Investigator: ScreenConnect phishing cluster",
"query": "<cenql_from_report>"}'
```

Before submitting a bulk IR investigation, check remaining weekly allowance to avoid exhausting the quota mid-incident.

```
$ curl -X GET 'https://api.platform.censys.io/usage?org_id=
<org_id>' \
  -H 'Authorization: Bearer <token>'
```

— Names endpoints, params, and quotas with runnable curl examples.

launch-20261027-3a04b411

02 **CenQL clusters saved as watched Collections** NEW 69

Investigation results, specifically the CenQL cluster query returned by Investigator, can be piped directly into `POST /v3/collections` to create a persistent, auto-updating watched cluster, removing the need to manually re-run hunts.

Upload a vendor threat-intel report and expand its indicators into a full campaign cluster, then save the CenQL query as a watched Collection.

```
$ # 1. Upload the report HTML as a file
curl -X POST 'https://api.platform.censys.io/files?org_id=
<org_id>' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"content_type": "text/html", "filename": "vendor-
report.html", "size_bytes": 42800}'

# 2. PUT the actual bytes to the returned upload_url, then run
investigation
#    (investigate() wraps the investigation submission call with
file_ids=["<file_id>"], start_time="...")

# 3. Turn the returned CenQL into a persistent Collection
curl -X POST 'https://api.platform.censys.io/v3/collections?
org_id=<org_id>' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"name": "Investigator: ScreenConnect phishing cluster",
"query": "<cenql_from_report>"}'
```

— Names the chaining endpoint but mechanism of auto-update unexplained.

launch-20261027-3a04b411

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Munki

Feedback ⓘ v7.4.0 NOTES ↗

Managed software installation for macOS —

Munki 7.4.0 adds low-data deferral, optional uninstalls, relaunch arguments, custom button labels, `--check-hashes` for `makecatalogs`, and Linux build support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v7.4.0 https://github.com/munki/munki.git
# already have the repo? check out this version:
$ git checkout v7.4.0
```

Adds `--check-hashes` option to `makecatalogs` to optionally verify installer item hashes at catalog-build time.

Adds `RestartActionForUninstall` pkginfo key support across `managedsoftwareupdate`, Managed Software Center, `munkiimport`, and `makepkginfo` to specify a restart action specific to uninstalling an item.

Adds support for deferring large downloads when a low data connection is detected in `managedsoftwareupdate` and Managed Software Center.

Adds support for `optional_uninstalls` (Suggested Removals) in `managedsoftwareupdate` and Managed Software Center.

Adds support for additional launch arguments when Managed Software Center relaunches an application after an update.

+2 more

UPDATED SUPPORTED PLATFORM macOS support range updated

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

gsa

Feedback ⓘ v28.7.0 NOTES ↗ RANK

Greenbone Security Assistant - The web frontend for the Greenbone Community Edition

> GSA v28.7.0 ships asynchronous report export with a new activity manager, GMP-level export commands, multi-cron scanner scheduling with jitter control, and anonymized agent support bundle downloads.

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

01 Async report export with activity manager NEW

75

Adds `ReportExport` and `ReportsExport` GMP commands for exporting reports via the GMP protocol, and a new activity manager UI panel for tracking async report export jobs and downloading results once complete, without blocking the interface.

Track a long-running report export without blocking the UI — open the activity manager to monitor async export progress.

> 📍 In the GSA web UI, trigger a report export. Go to the Activity Manager panel (top navigation) to monitor the export job status and download the result when complete.

— Names both GMP commands and the UI panel, mechanism partly described. v28.7.0

02 Multiple cron schedules and max delay for scanners NEW

60

Adds support for multiple cron schedules in scanner settings, and a max delay field for agents scheduler cron times to control scheduling jitter.

Configure multiple cron schedules with a max delay for an agent-based scanner to spread scan load across time windows.

> 📍 In the GSA web UI, go to Scanners › (select scanner) › Edit › Scheduler Settings. Add multiple cron schedule entries and set an Agents max delay value for each cron time to control jitter.

— UI navigation given but no limits or defaults specified. v28.7.0

THINNER COVERAGE BELOW

03 Anonymized agent support bundle downloads NEW

50

Adds an anonymization option when downloading agent support bundles, letting users remove sensitive data before sharing bundles with support teams.

Download an agent support bundle with sensitive data removed for safe sharing with support teams.

> 📍 In the GSA web UI, navigate to Agents › (select agent) › Download Support Bundle and enable the anonymization option before downloading.

— Clear UI path but no detail on what gets anonymized. v28.7.0

04 **Scan type and status visibility in reports and tasks** IMPROVED

45

Shows scan type and status in the report details header, the report details information tab, and the task details header.

— Names exact UI locations but no mechanism or screenshot. v28.7.0

05 **Optional reportConfigId on report download** IMPROVED

45

Makes `reportConfigId` optional when downloading a report.

— Names the exact parameter but gives no context on prior behavior. v28.7.0

WAS THIS USEFUL? 👍 👎

gvmd

Feedback ⓘ v26.43.0 NOTES ↗ RANK

Greenbone Vulnerability Manager - The database backend for the Greenbone Community Edition

> gvmd v26.43.0 introduces Active Directory credential verification via GMP, adds encryption and anonymization for agent bundles, and extends the database schema with agent OS metadata and scheduler delay settings.

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

01 **Agent OS metadata and scheduler delay fields** NEW 55

Adds agent OS metadata fields and `max_delay_seconds` scheduler delay support for agent groups to the database schema.

— Names schema field but lacks usage detail v26.43.0

02 **Active Directory credential integration in GMP** NEW 50

Adds a new GMP command to verify AD credentials against an Active Directory integration, extends the credentials model with a new LDAP/AD credential type, and adds build and runtime feature flags for AD integration.

— No exact command name or flag values given v26.43.0

03 **Encryption and anonymization for agent bundles** NEW 25

gvmd now supports encryption and anonymization options for agent bundles.

— No mechanism or option names given v26.43.0

WAS THIS USEFUL? 👍 👎

Rapid7 Insight Platform

Feedback ⓘ

changelog · Oct 27, 2026 | BLOG ↗ | RANK

Rapid7 Insight Platform is a cloud-based cybersecurity suite that unifies vulnerability management, SIEM, threat detection, and incident response across an organization's IT environment.

> Rapid7 Command Platform adds SCIM support for automated user provisioning and deprovisioning through identity providers such as Okta.

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

01 SCIM-based user provisioning in Command Platform NEW

58

Command Platform now supports SCIM for automated onboarding and offboarding, so accounts are created and deactivated in sync with changes in an identity provider such as Okta. SCIM is disabled by default and must be enabled by a platform administrator through Command Platform's SSO Settings, where the SCIM section lives alongside the identity provider configuration steps.

Enable SCIM to automate employee onboarding and offboarding so accounts are created and deactivated in Command Platform when changes occur in your IdP.

> 📍 In Command Platform, go to SSO Settings, locate the SCIM section, and enable SCIM integration — then follow the identity provider configuration steps (e.g. in Okta) to connect your IdP.

— Names UI path and IdP example but no API/endpoint detail 2026-10-27-39e127b7

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Privileged Access Management

Palo Alto Networks CyberArk

Feedback ⓘ

announced · Oct 28, 2026 BLOG ↗

Cyberark is a protect your organization with Idira®. Modernize identity governance, automate access reviews, and secure every identity across your hybrid cloud environment.

Idira Agent Identity Security is now GA with Prisma AIRS, delivering per-agent SPIFFE/OAuth 2.1 credentials, MCP tool-level ZSP, on-behalf-of audit logging, shadow agent onboarding, and an activity map.

NEW SUPPORTED PLATFORM Provisions SPIFFE Verifiable Identity Documents (SVIDs) and OAuth 2.1 client credentials for each registered agent, issuing short-lived, time-bound access tokens in place of long-lived shared secrets.

Applies fine-grained access policies to specific tools available on an MCP server, enabling least-privilege and zero standing privilege (ZSP) controls per agent.

NEW SUPPORTED PLATFORM Delivers ZSP for databases via the Secure Infrastructure Access MCP server, with ZSP for cloud providers listed as coming soon.

NEW INTEGRATION Supports automatic OAuth app registration for MCP servers that implement Dynamic Client Registration (DCR), capturing audit data for tool interactions.

Brings shadow AI agents into central governance through an unmanaged agent onboarding flow, including agents identified by cloud-provider origin IDs such as AWS ARNs or App IDs.

+4 more

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

Okta for AI Agents

Feedback ⓘ changelog · Oct 28, 2026

CHANGELOG ↗ RANK

Okta for AI Agents governs AI agent identities, permissions, and access across autonomous workflows.

> Okta adds mTLS client authentication for OIDC apps, replaces the flat authenticator enrollment model with configurable enrollment sets, and extends device visibility with a new Device Bound SSO card for managed macOS — alongside a wave of new OIN integrations and a Cross App Access breaking change tied to Managed Connection removal.

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

01 **Device Bound SSO card for managed macOS** NEW 75

Adds a Device Bound SSO (DBSSO) card to the device details page for managed macOS devices, showing DBSSO session status per OS-level user account alongside existing Platform SSO and Okta FastPass enrollment status.

Verify that a managed macOS device's users have active device-bound sessions without switching between multiple admin screens.

> 📍 In the Admin Console, go to Devices, select a managed macOS device, and open the device details page. The new Device Bound SSO card shows the DBSSO session status for each OS-level user account on that device alongside their Platform SSO and Okta FastPass enrollment status.

— Exact navigation path and mechanism both given changelog-20261028-c7c6e5d6

02 **New OIN integrations across provisioning, SSO and SCIM** NEW 65

Adds new OIN integrations: Sumo Logic provisioning (by Actions, covering user create, update, deactivate/delete, and role assignment across Sumo Logic regional deployments), Cloudflare provisioning, LeakData Shared Signals (API Service Integration), Ollet (IDV), Ready Room (OIDC), Crossbeam (SAML), Guru (SAML), Guru (SCIM), and Brava Security (SCIM).

— All integrations named with type but no setup steps changelog-20261028-c7c6e5d6

03 **Managed Connection removed from Cross App Access** BREAKING 65

The Managed connection tab on the app instance page in the Admin Console has been removed. Existing Cross App Access configurations using Managed Connection will stop working soon — admins must reconfigure connections using Cross App Access before the cutover to avoid disruption.

— Exact UI location named but no migration steps detailed changelog-20261028-c7c6e5d6

04 **mTLS client authentication for OIDC apps** NEW 60

Adds mTLS client authentication for OIDC apps, issuing access tokens cryptographically bound to a client certificate so stolen tokens cannot be reused without it — eliminating the need to store, share, or rotate shared client secrets.

— Mechanism explained but no config key or UI path given changelog-20261028-c7c6e5d6

05 **Authenticator enrollment sets** NEW 60

Adds authenticator enrollment sets, letting admins define groups of authenticators and require a subset (e.g. two of three) or let end users choose which authenticator to enroll, replacing the flat required/optional/disabled model.

— Behavior change described but no setup steps or API given changelog-20261028-c7c6e5d6

THINNER COVERAGE BELOW

06 **MCP server registration UI aligned with agent management**

IMPROVED 40

Updates the MCP server registration UI to match the AI agent and Agent Gateway management experiences, providing a consistent workflow for registering and configuring MCP servers.

— UI consistency noted but no concrete fields or steps changelog-20261028-c7c6e5d6

└──▷ **BREAKING ON UPGRADE**

! The Managed connection tab on the app instance page in the Admin Console has been removed. Existing Cross App Access configurations using Managed Connection will stop working soon — reconfigure connections using Cross App Access before the cutover to avoid disruptions.

WAS THIS USEFUL? 👍 👎

Aembit

Feedback ⓘ

Aembit Now Supports SPIFFE: X.509-SVID and JWT-SVID Without SPIRE (2026-10-27)

BLOG ↗ RANK

Aembit is a runtime IAM for AI agents and workloads. Enforce access policies, keep downstream credentials out of agents, audit access, and revoke it with a kill switch.

> Aembit added native SPIFFE support, issuing both X.509-SVID certificates and JWT-SVID tokens directly from its control plane without requiring a SPIRE server or agent.

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

01 **X.509-SVID certificates for workload mTLS** NEW 78

Aembit issues SPIFFE-compliant `X.509-SVID` certificates for end-to-end mTLS between workloads, delivered via Agent Proxy, CLI, or Aembit Secrets Operator (Kubernetes). Certificates auto-rotate at 80% of their lifetime (15-minute default, configurable), and private keys are never persisted to disk — all issued directly from the Aembit control plane without needing a SPIRE server or agent.

— Rotation threshold, default lifetime and delivery paths named, no sample command

Aembit Now Supports SPIFFE: X.509-SVID …

02 **JWT-SVID tokens for bearer-token auth** NEW 71

Aembit issues SPIFFE `JWT-SVID` tokens for bearer-token authentication behind service meshes such as Istio and Envoy, automatically populating standard claims (`sub` , `iss` , `aud` , `exp` , `iat` , `jti`) and supporting custom claims built from literal values or dynamic workload attributes like Kubernetes namespace, AWS account ID, and environment.

— Claims and mesh integrations named, no config syntax shown

Aembit Now Supports SPIFFE: X.509-SVID …

WAS THIS USEFUL? 👍 👎

Token Security

Feedback ⓘ announced · Oct 26, 2026 BLOG ↗

RANK

Token Security accelerates secure adoption of Agentic AI by managing & governing AI agents & non-human identities.

> Token Security launches Rehearsal, an AI-driven walkthrough mode in Token AI that navigates the platform, applies filters and answers plain-language questions on a user's behalf.

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

01 **Rehearsal AI-driven walkthrough in Token AI** NEW 77

Rehearsal is an AI-driven walkthrough mode accessible via Token AI that accepts plain-language goals, plans a route through the platform, opens pages, sets filters, highlights results, and writes the answer — all within the user's own browser tab. It runs as read-only automation (no secret fields are filled and no blocked calls are sent) and can be paused at any time by clicking to take back control, then resumed from where it left off.

Discover which AI agents have no owner without manually navigating filters — type a plain-language goal into Token AI and let Rehearsal walk the route for you.

> 📍 Open Token AI › type: 'Which AI agents have no owner?' › Rehearsal opens the relevant page, sets filters, and highlights matching agents.

Bootstrap a new playbook without knowing the exact UI path — ask Rehearsal to create it and it will navigate to the Playbooks page and scaffold the workflow.

> 📍 Open Token AI › type: 'Create a playbook that sends a Slack message when an AI agent gets a critical finding' › Rehearsal navigates to Playbooks, drafts the configuration, and prompts before saving.

— Mechanism and constraints described, with concrete example prompts to try.

launch-20261026-344bb54e

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Passbolt

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗ RANK

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

> Passbolt 5.17.0 adds phishing-resistant passkey MFA to its browser extension and strengthens several existing browser extension behaviors, including permission confirmations, autofill detection, and Manifest V3 support for Microsoft Edge.

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

01 **Edge extension moved to Manifest V3** IMPROVED 50

Moves the Microsoft Edge extension to Manifest V3 for continued credential access as Microsoft phases out Manifest V2.

– Names platform and manifest version but no migration steps. changelog-20261028-d0791bcc

02 **Passkey MFA in browser extension** NEW 45

Adds phishing-resistant passkey multi-factor authentication to the browser extension.

– States the capability but no setup steps or scope details. changelog-20261028-d0791bcc

03 **Permission change confirmation on resource moves** IMPROVED 35

The browser extension now confirms permission changes when resources are moved.

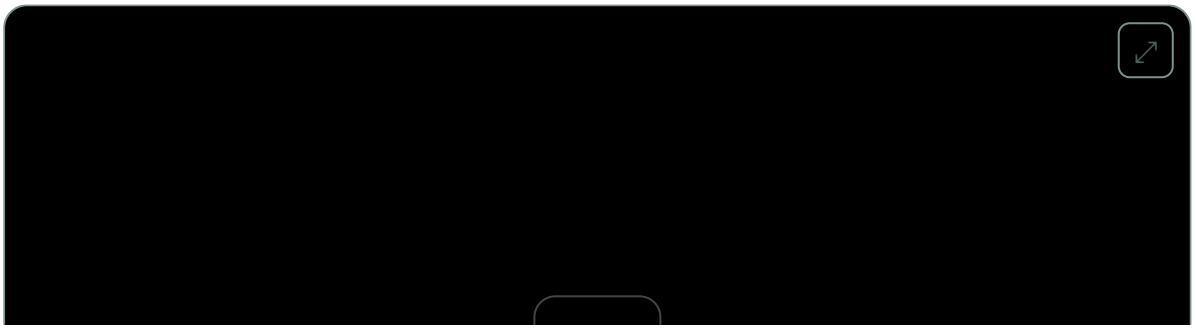
– Names the behavior but not the exact UI flow. changelog-20261028-d0791bcc

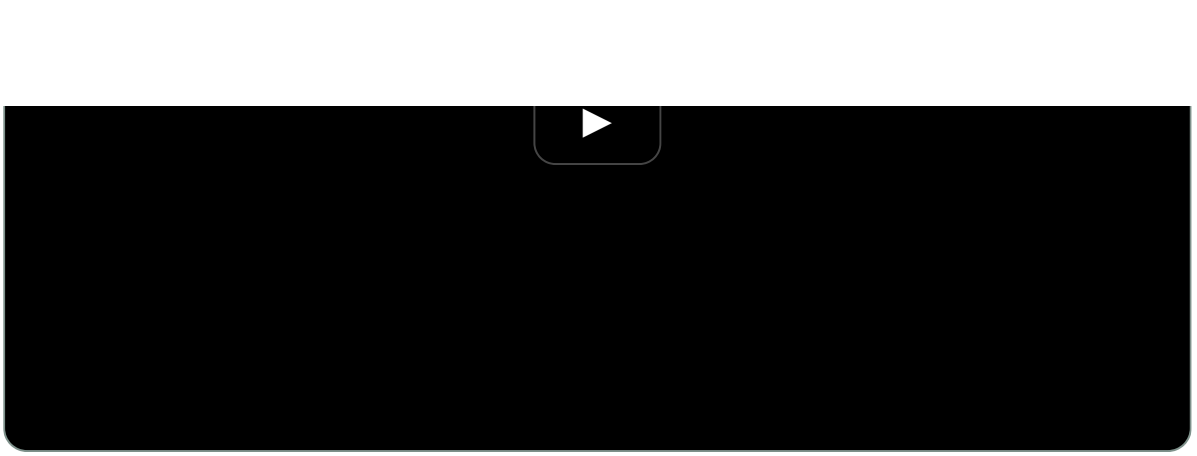
04 **Improved autofill detection** IMPROVED 30

Improves autofill detection across modern websites in the browser extension.

– Vague on which sites or mechanisms changed. changelog-20261028-d0791bcc

└──▷ **ALSO FROM THESE RELEASES**





WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗ RANK

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

> Auth0 shipped a major CLI expansion for AI agent workflows alongside early-access releases of CIMD dynamic client registration, Application Profiles governance, a new Brute-Force Protection unblock mode, and graduation of the FGA Permissions Index to Early Access.

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

01 **Agent Mode and AI-agent tooling in Auth0 CLI** NEW 95

The CLI adds `--agent-mode` (or `AUTH0_AGENT_MODE=true` env var) for structured JSON output, prompt-free operation, guaranteed non-zero exit codes on failure, and stable error `code` / `reason` fields, auto-detected for recognized agent clients. It also adds `auth0 commands --flat` (with optional `--detailed`) for full CLI discovery in one call, `--schema` to print the exact JSON payload a command expects (sourced live from the OpenAPI spec), `--data` to send a whole payload as one JSON object with local validation before it reaches the API, `auth0 docs search` for searching Auth0 documentation from the terminal without login, and `auth0 agent skills install` to install the Auth0 skill into Claude Code, Cursor, or all supported assistants.

— Names every flag and subcommand with exact syntax changelog-20261028-e47924d4

02 **Manual Unblock Behavior for Brute-Force Protection** NEW 80

An Unblock Behavior setting has been added to Brute-Force Protection (`Auth0 Dashboard > Security > Attack Protection > Brute-Force Protection`), with Automatic (default) and Manual modes; Manual mode prevents email security scanners from consuming recovery links before the user opens them.

Configure Manual unblock mode so email security scanners cannot accidentally consume recovery links and unblock accounts prematurely.

📍 Auth0 Dashboard › Security › Attack Protection › Brute-Force Protection › Unblock Behavior › select Manual › Save

— Exact dashboard path and setting name with runnable steps changelog-20261028-e47924d4

03 **Application Profiles for CIMD governance** NEW 65

Application Profiles, in Early Access, enable centralized governance of security settings — refresh token lifetimes, leeway, DPoP, and B2B login behavior — across all CIMD-

registered applications from a single configuration point in the Auth0 Dashboard under the Applications section.

— Names dashboard location and governed settings, no exact field names

changelog-20261028-e47924d4

04 **FGA Permissions Index reaches Early Access** IMPROVED

65

The FGA Permissions Index graduates from Developer Preview to Early Access, with self-service build/monitor/manage inside the FGA Dashboard. Enterprise subscriptions include one index with up to 10 million expansions, with a paid add-on scaling to 10 indexes and higher expansion limits.

— Concrete limits given but no config/API specifics changelog-20261028-e47924d4

THINNER COVERAGE BELOW

05 **CIMD dynamic client registration** NEW

50

Client ID Metadata Document (CIMD) with dynamic registration, in Early Access, lets MCP clients and AI agents connect to a tenant via just-in-time registration using a metadata URL, with no pre-registration required.

— Explains mechanism but no concrete config steps given changelog-20261028-e47924d4

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL?  

iDenfy

Feedback ⓘ snapshot-20261028 NOTES ↗ RANK

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

> iDenfy adds a dashboard email delivery log, lets users edit or delete their own comments on KYB/AML records (with deleted comments now hidden from the API), and gives unverified accounts five free KYB checks.

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

01 Comment editing and deletion on KYB/AML records

BREAKING

63

Comment authors can now edit or delete their own comments on KYB companies, AML v3 checks, monitoring records, and risk assessments. As a consequence, deleted comments are no longer returned by the API.

– Names affected record types and API behavior change snapshot-20261028

THINNER COVERAGE BELOW

02 Email delivery log in dashboard NEW

55

A new log in the dashboard records every email the platform has sent, showing the delivery channel and send status for each one. Reviewable via `Logs > Email Log`, where each row lists the recipient, delivery channel, and send status.

Review email delivery status for a specific outbound notification directly from the iDenfy dashboard.

> 📍 In the iDenfy dashboard, navigate to Logs › Email Log. Each row shows the recipient, delivery channel, and send status for every email the platform has sent on your behalf.

– Names exact UI navigation path and fields shown snapshot-20261028

03 Free KYB checks for unverified accounts NEW

35

Accounts without business verification can now run 5 free KYB checks.

– States the quota but no further mechanism or limits snapshot-20261028

04 Predictable Soft KYC results for dummy people IMPROVED

25

Soft KYC now returns predictable results for dummy people in the development environment.

— Brief statement with no mechanism or example given `snapshot-20261028`

└──▷ **BREAKING ON UPGRADE**

! Deleted comments are no longer returned by the API — callers that relied on retrieving deleted comment records will receive empty or absent entries.

WAS THIS USEFUL? 👍 👎

Castle

Feedback ⓘ announced · Oct 27, 2026 BLOG ↗ RANK

Castle is a complete suite of SDKs and APIs to monitor users, block bots, and stop in-app fraud.

> Castle ships AI agent identification, exposing a structured `agent` object in its API response and giving teams dashboard visibility and policy controls over AI agents and bots.

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

01 **AI agent identification via `agent` object** NEW 90

Castle adds a top-level `agent` object to the Castle API response with fields `name`, `operator`, `category`, and `identified_by`. `category` supports `ai_agent`, `ai_browser`, and `browser_automation`; `identified_by` supports `web_bot_auth`, `fingerprint`, and `user_agent`. It identifies named agents and browsers — including Meta Muse, OpenAI Dots, Grok Bot, ChatGPT agent, Instinct, and Perplexity Comet — using Web Bot Auth cryptographic signing, fingerprinting, and user-agent string analysis. Teams can filter and investigate events by agent identity in Castle's event dashboard, and set agent-specific policies to allow, restrict, or require additional verification for specific agents or agent categories based on the operation performed.

Inspect which AI agent triggered a Castle event — and how it was identified — to decide whether to allow, challenge, or block the request.

```
$ {
  "agent": {
    "name": "ChatGPT agent",
    "operator": "openai",
    "category": "ai_agent",
    "identified_by": ["web_bot_auth"]
  }
}
```

— Names exact API fields and values plus dashboard/policy use, but UI path not given.

launch-20261027-14cd00d2

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ

Stop nudging deleted, archived, and suspended users and accounts NOTES ↗

RANK

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

> Nudge Security now prevents nudges from being sent to deleted, archived, or suspended users and accounts, with clearer feedback when recipients are skipped.

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

01 Skip nudges to inactive users and accounts IMPROVED

45

The nudge button is disabled with an explanatory message when the target user or account is deleted, archived, or suspended. When sending bulk nudges to a mix of active and inactive recipients, Nudge Security now reports which recipients were skipped.

— Explains behavior but no config/flag named Stop nudging deleted, archived, and sus…

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ Campaign scope by users (2026-10-23)

BLOG ↗ RANK

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

> ConductorOne introduced a By users review scope for access review campaigns, letting reviews start from people rather than entitlements, with support for exclusions and automation triggers on role change or departure.

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

01 **'By users' review scope for access campaigns** NEW 72

Adds a new By users review scope type to access review campaigns, letting reviewers target specific users, user criteria, or a CEL expression instead of selecting individual apps and entitlements. Supports excluding entitlements and narrowing which accounts and grants are in scope within a By users campaign, and extends the scope to campaign templates so automations can trigger a review of a user's current access on role change or departure.

Trigger an access review automatically when a user changes roles or leaves, using a campaign template with By users scope paired with an automation.

📍 In the ConductorOne console, go to Access Reviews › Campaign Templates, open or create a template, set the Scope type to By users, configure user criteria or a CEL expression to match the departing/role-changing user, then go to Automations and create or edit an automation that starts this template when a user's role changes or their account is flagged for offboarding.

— Explains mechanism and scope but lacks named API/config surfaces

Campaign scope by users (2026-10-23)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password SaaS Manager #2231 (2026-10-21)

BLOG ↗

All the latest releases from the 1Password team

Adds budget assignment for people and roles; improves Access tab clarity with searchable access levels; integration results show AI usage record counts.

Adds budget assignment for people and application roles, with budget owners able to view budget details and receive AI usage and workflow notifications.

Updates the application Access tab to show access levels, provisioning methods, and access policies more clearly, with search by name or description, catalog settings view, and unsaved-change warnings.

Integration results now display the total number of AI usage records imported when AI usage is supported and enabled.

UPDATED INTEGRATION Atlassian Cloud connector update

UPDATED INTEGRATION Google Workspace connector update

UPDATED INTEGRATION LinkedIn Recruiter and LinkedIn Sales Navigator identity matching improvement

UPDATED INTEGRATION Connector-specific improvements for Paylocity, BambooHR with Pathify, Xero, and Marktplaats

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Zscaler Internet Access

Feedback ⓘ

changelog · Oct 27, 2026 | BLOG ↗ | RANK

Zscaler Internet Access is a cloud-based secure web gateway that routes user traffic through a zero-trust proxy to enforce security policies, block threats, and control access to the internet and SaaS applications.

> Zscaler Cloud Sandbox gains AI-driven inline verdicts, macOS detonation support, larger file and archive inspection limits, and an on-demand API analysis workflow, with an endpoint sandbox and content disarm & reconstruction planned for early 2027.

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

01 AI/ML detection visibility and inline instant verdicts NEW

70

Insights Logs now include an `AI/ML Detection Source` column for reviewing AI/ML-driven detections and actions inline. Files scoring 91–100 on the AI/ML threat confidence scale receive an AI Instant Verdict and are blocked inline before reaching an endpoint, with roughly 91% receiving instant benign verdicts and about 9% escalated to full dynamic detonation.

— Names log column and score thresholds but only a UI path to act on 2026-10-27-034134aa

02 API-Driven File Analysis NEW

65

Adds API-Driven File Analysis for on-demand file submission, producing MITRE ATT&CK-mapped reports with TTPs, IOCs, and PCAPs, feeding into SIEM/SOAR workflows.

— Describes output and integration but no endpoint or method named 2026-10-27-034134aa

03 Larger file and archive inspection limits IMPROVED

60

Raises file size support for EXE and DLL inspection to 200 MB, with 10-layer archive inspection for nested and compressed packages.

— Concrete size and depth limits but no activation steps 2026-10-27-034134aa

THINNER COVERAGE BELOW

04 macOS detonation support in Cloud Sandbox NEW

50

Cloud Sandbox adds macOS detonation support for `.dmg` and `.pkg` files across both inline and API channels, extending behavioral analysis to Mac fleets.

— Names file types and channels but no setup steps given 2026-10-27-034134aa

05 **Upcoming sandbox capabilities for early 2027** NEW 45

Plans an Endpoint Sandbox (coming early 2027) via Zscaler Client Connector to intercept and inspect files arriving via offline paths such as USB, Bluetooth, and AirDrop before execution, and Content Disarm & Reconstruction (coming early 2027) to strip malicious content, reconstruct files in their original native format, and deliver sanitized versions.

— Roadmap items not yet available, nothing actionable today 2026-10-27-034134aa

WAS THIS USEFUL? 👍 👎

Cato Networks

Feedback ⓘ

Cato AI Endpoint secures browser AI, desktop AI and coding agents on every endpoint (202

BLOG ↗ RANK

Cato networks is a cato SASE Cloud converges SD-WAN, security, global backbone, and remote access into a global cloud-native service. WE ARE SASE™

> Cato Networks launched Cato AI Endpoint, a new lightweight client that discovers and governs AI usage on endpoints — from browser AI and desktop AI apps to local coding agents and MCP servers — alongside existing EDR and SSE stacks.

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

01 **Cato AI Endpoint client for AI and coding agents** NEW

60

Cato AI Endpoint is a new lightweight client that secures browser AI, desktop AI apps, and local coding agents (Claude Code, Codex, Cursor) from a single dashboard, coexisting with existing EDR and SSE stacks. It discovers Shadow AI — unapproved AI applications, local agents, MCP servers, and skills installed on endpoints — to build a governed inventory before policy is written. Each skill gets a risk finding with a block, monitor, or allow recommendation, and verified remote MCP servers are badged so teams can prioritize the riskiest agents. Endpoint and network AI events are correlated in a shared platform and data lake, surfacing the user, endpoint, and data behind each blocked tool or prompt in one dashboard.

— Describes mechanism and scope well but no CLI, API or config surface named.

Cato AI Endpoint secures browser AI, de…

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ

Device Provisioning with OAuth Apps (2026-10-27) BLOG ↗ RANK

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

> Tailscale made OAuth Apps for Device Provisioning generally available via its API.

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

01 **OAuth Apps for Device Provisioning now GA** NEW 40

Tailscale OAuth Apps for Device Provisioning can now be created and managed via the Tailscale API, with the feature moving to general availability.

— Names the API surface but no endpoint or parameters given

Device Provisioning with OAuth Apps (20…

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ October 27, 2026 NOTES ↗

RANK

DigiCert CertCentral is a web-based platform for ordering, issuing, managing, and renewing SSL/TLS certificates and other digital trust credentials.

> DigiCert CertCentral Europe now automatically matches intermediate CA chain key types to CSR key types and warns on mismatches, giving practitioners more visibility and control when selecting ICA chains for dual RSA/ECC products.

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

01 ICA chain key-type matching and mismatch warnings NEW

79

In CertCentral Europe, for products with ICA selection enabled that support both RSA and ECC, the platform automatically selects an ICA chain whose key type matches the CSR key type when a product's default ICA chain is set to DigiCert default. If a specific ICA chain is selected whose key type differs from the CSR (e.g. an ECC certificate with an RSA ICA chain), a key-type mismatch warning appears on the request form via the `Intermediate chains [Intermediate CA] > [Root CA]` menu, but order submission with the mixed chain is still allowed. Administrators can check ICA selection status and the default chain per product under `Settings > Product Settings` using the Allowed intermediate chains and Default intermediate chain menus.

Verify whether ICA selection is enabled for a specific product and check its default chain before ordering, to know whether automatic key-type matching will apply.

📍 In CertCentral Europe, go to Settings › Product Settings. Select the product (e.g. Basic OV). Look for the 'Allowed intermediate chains [Intermediate CA] › [Root CA]' menu. If present, check the Default intermediate chain value: DigiCert default means auto key-type matching applies; a specific chain (e.g. 'DigiCert Global G3 TLS ECC SHA384 2020 CA1 (SHA384ECDSA) › DigiCert Global Root G3 (SHA384ECDSA)') means a mismatch warning may appear if your CSR key type differs.

— Mechanism and exact UI paths given, but no CLI/API surface. October 27, 2026

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Backup & Recovery

Veeam Backup & Replication

Feedback ⓘ

changelog · Oct 27, 2026 NOTES ↗ RANK

Veeam Backup & Replication is backup, recovery, and replication software for virtual, physical, cloud, and SaaS workloads.

> Veeam Backup & Replication 13.1.2.57 adds VMware vCenter Server 9.1.1 support and bundles the Google Cloud Plug-In v8.1, while enforcing a new 90-day minimum archive window for Veeam Data Cloud Vault Archive and automatically excluding certain NetApp ONTAP folders from file share backups.

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

01 **90-day minimum archive window for Data Cloud Vault** BREAKING

75

When Veeam Data Cloud Vault Archive is used as the archive tier of a scale-out backup repository, the archive window must now be at least 90 days. This is enforced consistently across the console, Web UI, and PowerShell, and existing configurations with a shorter window will be rejected.

— Clear mechanism and scope, no exact PowerShell cmdlet named 2026-10-27-be477f60

02 **NetApp ONTAP file share backup exclusions** IMPROVED

60

Backups of file shares from NetApp ONTAP now automatically exclude the `.copy_offload` and `.anti_ransomware_analytics_log` folders.

— Names exact folders excluded but no config toggle given 2026-10-27-be477f60

THINNER COVERAGE BELOW

03 **Google Cloud Plug-In v8.1 bundled in ISOs** NEW

50

`Veeam Plug-In for Google Cloud v8.1` is now bundled in both the Veeam Backup & Replication ISO and the Veeam Software Appliance ISO.

— Names exact plugin version and install media, no further detail 2026-10-27-be477f60

04 **VMware vCenter Server 9.1.1 support** NEW

45

Adds support for VMware vCenter Server 9.1.1, eliminating spurious warnings about untested API versions that previously appeared with this vCenter release.

— Explains the fix but no reproduction or config detail 2026-10-27-be477f60

└──▷ **BREAKING ON UPGRADE**

! When Veeam Data Cloud Vault Archive is the archive tier of a scale-out backup repository, the archive window must now be at least 90 days; existing configurations with a shorter window will be rejected by the console, Web UI, and PowerShell.

WAS THIS USEFUL?  

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗ RANK

RegScale is a compliance automation platform that helps organizations manage security controls, risk, and regulatory frameworks like NIST, FedRAMP, and SOC 2 through continuous monitoring.

> RegScale 6.55.0 adds per-change-type CCB routing for DevSecOps approvals, strict exit codes for failed imports and syncs, Python 3.15 support, uv-based packaging, and GCP SCC CVE ingestion, alongside a wide set of breaking fixes that tighten validation and failure reporting across AWS, GCP, and Wiz syncs.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 **CCB routing for DevSecOps change approvals** NEW 91

Adds `change.approval_by_type` config key and `--approval-mode` CLI flag to route DevSecOps release/deploy Change records to direct completion, policy approval, or CCB per change type, with a built-in default for Emergency Change records (`change.approval_by_type.Emergency: ccb`) where policy rejection is reported but does not fail the job. `devsecops change request-approval` now reports `--approval-mode` status, surfacing the Change status and policy mode returned by the platform and failing only on enforced rejections.

Route Emergency Change records to the CCB while sending standard release Changes directly, without requiring CCB review.

yaml

```yaml
change:
  approval_by_type:
    Emergency: ccb
    Release: direct
    Deploy: request
```

— Names config key, flag, default and command together changelog-20261028-d0356797

02 **AlienVault OTX server override** NEW 81

Adds an `otxUrl` setting (environment variable `REGSCALE_OTX_URL`) so `alienvault ingest_pulses` can point at a custom AlienVault OTX server instead of the default `https://otx.alienvault.com`.

Point AlienVault OTX ingestion at an on-premise OTX mirror instead of the public cloud endpoint.

```
$ export REGSCALE_OTX_URL=https://otx.internal.corp
regscale alienvault ingest_pulses
```

— Exact env var, default URL and command given changelog-20261028-d0356797

03 **GCP SCC findings without NIST mapping** NEW 78

`gcp sync_findings` now imports SCC findings with no NIST 800-53 control mapping, including every `SOFTWARE_VULNERABILITY` (CVE) finding, keyed on SCC finding name with the CVE id and CVSS v3 base score.

— Names command, finding type, keying fields changelog-20261028-d0356797

04 **AWS sync_cloudwatch control failure threshold** BREAKING 78

AWS `sync_cloudwatch` and `sync_cloudwatch_ex` now fail AU-6, AU-9, AU-11, AU-12, and SI-4 controls when more than 20% of log groups fail them; assessments that previously passed because one log group passed are rewritten as Fail on the first run after upgrading.

— Specific controls, threshold and upgrade impact named changelog-20261028-d0356797

05 **set_permissions load validation changes** BREAKING 77

`set_permissions load` now exits with an error (non-zero) after trying every row when the platform rejects a row or a row holds an invalid value, instead of logging a warning and exiting 0; it also now rejects rows where the `public | private` cell is blank or mistyped, instead of silently defaulting the record to public.

— Names command, field and before/after behavior changelog-20261028-d0356797

06 **Strict exit codes for JCAM and Security Hub failures** IMPROVED 73

`REGSCALE_STRICT_EXIT_CODES=1` now makes JCAM imports exit 3 when records could not be imported, and makes AWS Security Hub syncs exit 3 for assets and findings that could not be processed, instead of silently exiting 0 in both cases.

— Env var and exit codes named but no sample command changelog-20261028-d0356797

07 **Packaging migration from Poetry to uv** BREAKING 70

Migrates packaging and dependency management from Poetry (`poetry-core` , `poetry.lock`) to `uv` (`hatchling` , `uv.lock`) for faster dependency resolution and container builds; build pipelines using Poetry must be updated.

— Names old and new toolchain files, has migration note changelog-20261028-d0356797

08 **AWS sync_findings consolidated mode dedup** BREAKING

70

AWS `sync_findings` in consolidated mode replaces merged per-product catch-all vulnerability records with one record per Inspector CVE and one per GuardDuty finding type; the next unscoped sync closes the old catch-all records.

— Named command, products and record granularity change changelog-20261028-d0356797

09 **GCP sync_compliance auto-pass removal** BREAKING

65

`gcp sync_compliance` no longer marks a control Pass or Fully Implemented when no SCC finding maps to it; controls previously auto-passing with no evidence will now be left unevaluated or failed.

— Clear before/after but no exact command flags changelog-20261028-d0356797

10 **AWS command authentication failure exit code** BREAKING

60

Every `aws` command now exits 4 on authentication failure before reading or writing anything, instead of reporting completed successfully and syncing an empty inventory.

— Exact exit code and prior silent-failure behavior named changelog-20261028-d0356797

THINNER COVERAGE BELOW

11 **Wiz asset disk storage value change** BREAKING

57

Wiz `inventory` and `stig_checks` now set an asset's Disk Storage (GB) to 0 instead of the virtual machine's disk count; existing Wiz assets show 0 after their next sync.

— Clear before/after value change, no workaround given changelog-20261028-d0356797

12 **AWS sync_cloudwatch account_id filtering fix** IMPROVED

55

Fixes `--account_id` filtering in `aws sync_cloudwatch` so it retains log groups in the specified account, rather than collecting none.

— Named flag and command, thin mechanism changelog-20261028-d0356797

13 **import assessment_linker partial list handling** BREAKING

48

`regscale import assessment_linker` now exits 1 with an error when RegScale returns only a partial assessment list, instead of linking a partial set and exiting 0.

— Named command and exit code, thin on scope changelog-20261028-d0356797

14 **Python 3.15 support** NEW 30

Adds Python 3.15 support with updated CI test matrices and wheel bounds.

— Bare statement, no further mechanism given changelog-20261028-d0356797

15 **regscale-dev update_docs command removed** DEPRECATED 25

The `regscale-dev update_docs` command has been removed.

— Single-line removal with no replacement noted changelog-20261028-d0356797

└──▷ **BREAKING ON UPGRADE**

! `set_permissions load` now exits with an error (non-zero) after trying every row when the platform rejects a row or a row holds an invalid value; previously it logged a warning and exited 0.

! `set_permissions load` now rejects rows where the `public | private` cell is blank or mistyped — previously a blank visibility silently made the record public.

! AWS `sync_cloudwatch` and `sync_cloudwatch_ex` now fail AU-6, AU-9, AU-11, AU-12, and SI-4 controls when more than 20% of log groups fail them; assessments that previously passed because one log group passed are rewritten as Fail on the first run after upgrading.

! `gcp sync_compliance` no longer marks a control Pass or Fully Implemented when no SCC finding maps to it; controls previously auto-passing with no evidence will now be left unevaluated or failed.

! AWS `sync_findings` in consolidated mode replaces merged per-product catch-all vulnerability records with one record per Inspector CVE and one per GuardDuty finding type; the next unscoped sync closes the old catch-all records.

! Wiz `inventory` and `stig_checks` now set an asset's Disk Storage (GB) to 0 instead of the virtual machine's disk count; existing Wiz assets show 0 after their next sync.

! Every `aws` command now exits 4 on authentication failure before reading or writing anything, instead of reporting `completed successfully` and syncing an empty inventory.

! `regscale import assessment_linker` now exits 1 with an error when RegScale returns only a partial assessment list, instead of linking a partial set and exiting 0.

! The `regscale-dev update_docs` command has been removed.

! Packaging and dependency lock files have migrated from Poetry (`poetry-core` , `poetry.lock`) to `uv` (`hatchling` , `uv.lock`); build pipelines using Poetry must be updated.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.40.0 (2026-10-28) BLOG ↗ RANK

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

> Runlayer 1.40.0 adds connector-level rate-limit policies and artifact secret scanning, alongside a waterfall agent-run timeline, GPT-6.1 Terra model support, voice messaging over iMessage, and a breaking 7-day cap on Incident API and MCP tool time-range queries.

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

01 **Rate-limit policies for connector tool calls** NEW 80

Admins can cap tool-call frequency per connector — for example 100 calls per hour per user or 30 calls per minute per session — using Rate limit as a new policy action, with two built-in rate-limit templates and per-policy badge display. Refused calls are logged in the audit log as Tool Call Rate Limited.

Cap a noisy connector to 30 tool calls per minute per session to prevent runaway agents from exhausting API quotas.

> 📍 In the Runlayer console, go to the connector's Policies tab, create a new policy, set Action to Rate limit, choose the '30 calls per minute per session' template, then save. Calls over the limit are refused and appear in the audit log as Tool Call Rate Limited.

— Names policy action, templates, and audit log entry 1.40.0 (2026-10-28)

02 **7-day cap on Incident API time-range queries** BREAKING 75

Incident API and MCP tool requests for a time range longer than 7 days now return only the 7 days ending at the range's end date; callers must request longer ranges in 7-day steps.

— Clear migration note naming API and exact behavior change 1.40.0 (2026-10-28)

03 **Workspace-wide artifact auto-delete policies** NEW 70

New auto-delete policies in Settings → Artifacts apply workspace-wide, with per-folder overrides such as Keep forever for a specific folder. Idle artifacts are archived first and deleted 30 days later.

Set a workspace-wide artifact auto-delete rule with a Keep forever override for a sensitive reports folder, so idle drafts are cleaned up while compliance artifacts are

retained.

> 📍 In the Runlayer console, go to Settings → Artifacts, enable the workspace-wide auto-delete rule, then add a per-folder override for the Reports folder set to `Keep forever`, and save.

— Names config path, override, and exact 30-day timing `1.40.0 (2026-10-28)`

04 **Waterfall view for agent run timelines** NEW 65

A new `Timeline / Waterfall` switch in the agent run drawer charts each tool call as a bar at its real offset from run start, with subagent steps nested under their parent call.

— Names UI control and exact visualization mechanism `1.40.0 (2026-10-28)`

05 **Clio session-topic filter for Sessions page** NEW 65

Adds a Clio v1.0 session-topic filter, gated behind the `session-topics-clio` feature flag, listing 63 Clio subtopics — for example `Research and prepare for sales engagements` — alongside existing topics on the Sessions page.

— Names feature flag, subtopic count, and example topic `1.40.0 (2026-10-28)`

06 **MDM device blocking for deactivated users** NEW 60

Deactivated users' MDM-managed devices now block tool calls in `Protect` and `Enforce` modes via AI Watch hooks; devices Runlayer cannot match to a user receive the organization's `Enforce` tool-type blocklist.

— Names modes and hooks but no user-facing action `1.40.0 (2026-10-28)`

THINNER COVERAGE BELOW

07 **Automatic retry for parked session compaction** IMPROVED 50

Session compaction now auto-retries parked sessions after upgrades, at most once per day and up to three retries per session, eliminating most manual support resets.

— Gives exact retry limits but no user-facing action `1.40.0 (2026-10-28)`

08 **Artifact secret scanning on publish** NEW 45

Runlayer blocks artifact publishing when content contains an obvious API key, access token, or private key, and instructs the agent to redact and republish.

— Describes trigger and behavior but no config surface `1.40.0 (2026-10-28)`

09 **Organization-scoped Metrics and 90-day range limit** IMPROVED 45

Plugin and skill install counts on the Metrics page, and a connector's most-active-users list, are now scoped to organization-only activity, and date ranges older than 90 days are rejected.

— States new scope and limit without usage steps 1.40.0 (2026-10-28)

10 **GPT-6.1 Terra model support** NEW 40

GPT-6.1 Terra is added to the agent model picker for OpenAI providers, with reasoning, image input, and Spend pricing.

— Names model and picker but lacks further detail 1.40.0 (2026-10-28)

11 **Voice messages over iMessage** NEW 40

Users can send a voice message from Messages and the agent transcribes it and replies, using the workspace's OpenAI provider.

— Explains flow but no config or limits given 1.40.0 (2026-10-28)

└──▷ **BREAKING ON UPGRADE**

! Incident API and MCP tool requests for a time range longer than 7 days now return only the 7 days ending at the range's end date; callers must request longer ranges in 7-day steps.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Security Posture & Discovery

Zenity

Feedback ⓘ announced · Oct 28, 2026 BLOG ↗

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

Zenity Boundaries adds chained Boundaries, Dogwood-style temporal logic, taint propagation, agent steering, and Production History replay for AI agent enforcement.

Adds chained Boundaries, letting a Boundary rule reference an earlier Boundary so a Detect match in one step can trigger a hard block on a sensitive tool call in a later step of the same agent session.

Extends Boundaries with Dogwood-style temporal conditions, enabling policies that reason about the sequence and timing of agent tool calls — such as blocking a `GitPush` unless tests passed within the last 15 minutes with no subsequent failure.

Introduces steering as an alternative to blocking, letting a Boundary return guidance to an agent about why an action was stopped and what an acceptable path looks like, so the agent can reroute rather than search for a workaround.

Supports taint propagation across a session, attaching labels (e.g. untrusted-source, credential-access, confidential-data) that later Boundaries can inspect to enforce the Lethal Trifecta pattern.

Provides Production History replay, allowing a new Boundary to be tested against the last seven days of real agent activity before being promoted to enforcement mode.

+1 more

NEW INTEGRATION Inline enforcement on Claude Code, Cowork, Microsoft Copilot Studio, Azure Foundry

NEW INTEGRATION Detection-only support for ChatGPT and Gemini

WAS THIS USEFUL? 👍 👎

◆ AI/LLM Security

NVIDIA OpenShell

Feedback ⓘ v0.1.6 NOTES ↗ RANK

OpenShell runs autonomous AI agents in policy-controlled sandboxes with isolated filesystem, network, and credential access.

> OpenShell v0.1.6 adds sandbox suspend/resume, an experimental shared Kubernetes supervisor topology, OCSF audit logging for gateways, several Helm chart additions, and a run of breaking changes to gateway authentication and the sandbox-supervisor connection protocol.

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

01 **Shared Kubernetes supervisor topology** NEW 95

Adds `supervisor_topology = "gateway"` under `[openshell.drivers.kubernetes]` to enable an experimental shared supervisor that hosts many sandboxes in one process as a pool that can take over a failed supervisor's sandboxes, with `supervisor.replicas` to set pool capacity. Gateway TLS is required when this topology is enabled.

Enable the shared Kubernetes supervisor so many sandboxes share one process and failed supervisors are covered by pool replicas.

toml

```toml
[openshell.drivers.kubernetes]
supervisor_topology = "gateway"

[openshell.gateway]
# gateway TLS is required when supervisor_topology = "gateway"

# In the Helm values, set replica count:
# supervisor:
#   replicas: 3
```

— Full config example and failover mechanism explained v0.1.6

02 **Gateways require authentication by default** BREAKING 85

Gateways now refuse to start unless OIDC, mTLS user or operator authentication, or gateway JWTs are configured; for trusted local development only, set `allow_unauthenticated_users = true` under `[openshell.gateway.auth]`.

— Clear before/after behavior and exact escape-hatch config key v0.1.6

03 **Prerelease snap beta channel** NEW 75

Prerelease builds now publish to the snap's `latest/beta` channel; running `OPENSHELL_VERSION=pre install.sh` installs from there.

Install the latest prerelease build via the snap beta channel to test upcoming features before a stable release.

```
$ OPENSHELL_VERSION=pre curl -LsSf
https://raw.githubusercontent.com/NVIDIA/OpenShell/main/install.s
h | sh
```

— Runnable install command with exact env var v0.1.6

04 **Sandbox suspend and resume commands** NEW 70

Adds `openshell sandbox suspend` and `openshell sandbox resume` commands to pause and resume sandboxes on supported compute drivers, freeing compute resources without losing state.

Pause a long-running sandbox to free compute resources, then resume it later without losing state.

```
$ openshell sandbox suspend my-sandbox
# ... later ...
openshell sandbox resume my-sandbox
```

— Named commands with a runnable example, but scope/mechanism thin v0.1.6

05 **Helm chart additions for env vars and PKI job metadata** NEW 70

Adds `server.extraEnv` to the Helm chart for passing extra environment variables, including Secret references, to the gateway container, and adds `pkiInitJob.podLabels` and `pkiInitJob.podAnnotations` for the certificate hook pods.

— Three named Helm keys with purpose, no usage example v0.1.6

06 **OCSF audit events for gateway activity** NEW 65

Gateways now write OCSF audit events for authorization denials, accepted credentials, SSH session tokens, and sandbox token checks, and raise a High-severity finding when a sandbox credential looks misused.

— Clear scope of events logged but no config or endpoint given v0.1.6

07 **Sandbox-to-supervisor connection direction reversed**

BREAKING

65

New sandboxes on Kubernetes, Docker, Podman, and VM drivers now open the control connection from sandbox to supervisor, the opposite direction from before; supervisor and sandbox runtime binaries must be upgraded together because older binaries reject the new sandboxes.

– Explains mechanism and upgrade requirement but no commands v0.1.6

08 **Middleware binding policy enforcement** NEW

60

Adds a `binding_policy` setting (values `allowlist` or `exact`) to require a middleware service to provide specific bindings such as request inspection.

– Named config key but no example or navigation path v0.1.6

THINNER COVERAGE BELOW

09 **ResumeBoundaryRequest field renamed** BREAKING

50

In external isolation backends, `ResumeBoundaryRequest` field 3 is renamed from `allow_cold_start_fallback` to `restart_if_restore_fails`.

– Exact field names given but minimal context or migration steps v0.1.6

10 **Versioned Go SDK module tags** IMPROVED

40

Stable releases now tag the Go SDK module as `sdk/go/vX.Y.Z` so Go tooling and `pkg.go.dev` show real versions.

– Names the tag format but gives no command to use it v0.1.6

└──▷ **BREAKING ON UPGRADE**

! Gateways now refuse to start unless OIDC, mTLS user or operator authentication, or gateway JWTs are configured; for trusted local development only, set `allow_unauthenticated_users = true` under `[openshell.gateway.auth]`.

! New sandboxes on Kubernetes, Docker, Podman, and VM now open the control connection from sandbox to supervisor in the opposite direction; supervisor and sandbox runtime binaries must be upgraded together because older binaries reject the new sandboxes.

! External isolation backends: `ResumeBoundaryRequest` field 3 is renamed from `allow_cold_start_fallback` to `restart_if_restore_fails`.

WAS THIS USEFUL?  

ToolHive

Feedback ⓘ v0.58.0 NOTES ↗ RANK

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

> ToolHive v0.58.0 adds an opt-in standards-compliant JWT type for embedded Authorization Server tokens and closes a gap allowing refresh grants on dead credentials.

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

01 **Opt-in at+jwt token type for embedded AS** NEW 50

ToolHive adds an opt-in `at+jwt` `typ` header for embedded Authorization Server access tokens, enabling standards-compliant JWT typing for downstream consumers.

— Names exact header value but no config/flag to enable it v0.58.0

02 **Refresh grants blocked for expired or revoked credentials**

IMPROVED 40

ToolHive now refuses refresh grants whose upstream credential has expired or been revoked, preventing silent use of dead credentials to obtain new tokens.

— Explains behaviour change but no named endpoint or config v0.58.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗ RANK

Obsidian is a note-taking app that lets you write and organize Markdown files stored locally in a vault.

> Obsidian 1.15.0 replaces its built-in Mermaid diagram renderer with a new library called Pele, adds voice capture with real-time transcription and PDF export on mobile, introduces heading-targeted URI/CLI commands, and ships smaller improvements to Bases, Vim mode, and search.

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

01 **Mermaid diagrams now rendered via Pele** BREAKING 85

Obsidian replaces the Mermaid library with a new internal renderer called Pele, compatible with Mermaid 12.1.0 syntax, with theme-matched colors/fonts, responsive sizing, and zoom/pan when diagrams are scaled down. It adds `--mermaid-*` CSS variables for colors, corner radius, and font weights, plus `.pele-*` classes on diagram elements for custom styling, and a loadMermaid() / `window.mermaid` compatibility shim supporting initialize(), render(), parse(), mermaidAPI.getConfig(), and setConfig() (new code can use `window.pele` once loadMermaid() resolves). It also adds loadPixi() to load the `PIXI` global before the first graph opens, since `PIXI` is no longer loaded at startup. Breaking: the Show Mermaid diagrams in notes setting is removed because Pele cannot execute code so diagrams always render; SVG output now carries Pele's `.pele-*` classes instead of Mermaid's classes; and Obsidian no longer bundles Mermaid.

— Mechanism and breaking changes fully named, but no runnable example

changelog-20261028-3b73e656

02 **append/prepend URI and CLI heading commands** NEW 75

Adds `append` and `prepend` URI/CLI commands to add content under a specific heading, with control over line breaks around the added content.

— Exact commands named and directly usable changelog-20261028-3b73e656

03 **Vim mode command and mapping additions** IMPROVED 75

Vim mode adds `g?`, `:put`, search patterns in `:` ranges (e.g. `:/foo/,/bar/d`), Vim-style regexes with `:set nopcre`, and `<A-…>` mappings on macOS.

— Every command and setting named verbatim and directly runnable

changelog-20261028-3b73e656

04 **Bases filter management and display improvements** IMPROVED 65

Bases gains dragging to reorder filters and a per-filter-row menu to add filters or move them between `All views` and `This view`, shows a result-limit notice (e.g. `Showing 10 of 40`) when search matches are hidden, and allows Formula fields to be resized up to the width of the window.

— Names all four UI additions but gives no exact navigation changelog-20261028-3b73e656

---

05 **Quick Capture via Spotlight and Siri** NEW 60

Adds Quick Capture as a Spotlight action and Siri shortcut (e.g. `Open Quick Capture in Obsidian`); setting Note Title to `Automatic` uses the first 80 characters of the capture as the title.

— Names shortcut phrase and setting but no exact navigation path

changelog-20261028-3b73e656

---

06 **Closing formatting token classes in Live Preview** NEW 60

Adds a `cm-formatting-*-end` class (e.g. `cm-formatting-strong-end`) to closing formatting tokens in Live Preview, matching the existing `cm-formatting-link-end` and `cm-formatting-highlight-end` classes.

— Names exact CSS classes usable for custom styling changelog-20261028-3b73e656

---

07 **Accent-insensitive search and suggestions** IMPROVED 60

Suggestions and search now ignore diacritics/accents, so for example `cafe` matches `café` in link suggestions, Quick Switcher, Command Palette, and search, and `prepareFuzzySearch` and `prepareSimpleSearch` now ignore accents on Latin and Greek letters.

— Names affected surfaces and API functions but no config toggle

changelog-20261028-3b73e656

THINNER COVERAGE BELOW

---

08 **Voice capture and transcription on iOS** NEW 45

Supports voice capture and real-time transcription on iOS 26, including from the Lock Screen, with recordings saved directly to the vault.

— Describes capability but no setting or flag named changelog-20261028-3b73e656

---

09 **PDF export for notes on mobile** NEW 30

Adds PDF export for notes on mobile.

— Bare feature name with no further detail changelog-20261028-3b73e656

└──▷ BREAKING ON UPGRADE

! The `Show Mermaid diagrams in notes` setting has been removed — diagrams always render now.

! loadMermaid() and `window.mermaid` now draw with Pele, not the Mermaid library. SVG output carries Pele's `.pele-*` classes instead of Mermaid's classes. Obsidian no longer bundles Mermaid.

! The `PIXI` global is no longer loaded at start-up; it is loaded the first time a graph opens. Use loadPixi() to load it earlier.

WAS THIS USEFUL?  

Ghostty

Feedback ⓘ 1.4.0 NOTES ↗ RANK

Ghostty is a fast, feature-rich, and cross-platform terminal emulator that uses platform-native UI and GPU acceleration.

> Ghostty 1.4 ships Kitty clipboard protocol support with a configurable write-size cap, transparent scrollback compression that lets the default history size triple, full Kitty graphics spec coverage, a new OpenGL rendering pipeline for Linux/FreeBSD, GTK split drag-and-drop, and macOS 27 support.

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

01 Kitty clipboard protocol and write-size limit NEW

91

Adds the Kitty clipboard protocol (OSC 5522) on macOS and GTK, letting terminal programs read and write arbitrary clipboard formats (images, HTML) and select among available formats; paste events via mode 5522 let programs detect what formats a paste contains (e.g. accepting pasted images). A new `clipboard-write-limit-bytes` configuration key (default 64 MiB) caps the size of a single clipboard write, with over-limit writes failing and leaving the clipboard untouched.

Cap clipboard writes to prevent a terminal program from writing an unexpectedly large payload to the clipboard.

```ini
clipboard-write-limit-bytes = 10485760
```

— Names exact config key, default, and runnable example 1.4.0

02 Scrollback compression and larger default history IMPROVED

75

Non-visible scrollback is transparently compressed while the terminal is idle and decompressed on demand, cutting physical memory usage by 70–90% for typical text-heavy history. This enabled raising the default scrollback limit from 10 MB to 50 MB per terminal.

— Clear mechanism and numbers but no config key given 1.4.0

03 New OpenGL renderer for Linux and FreeBSD NEW

73

Introduces a new OpenGL rendering pipeline on Linux and FreeBSD using EGL on its own thread, exporting frames as DMABUFs with triple-buffering, which eliminates blurry text under fractional display scaling.

— Detailed mechanism but nothing for user to configure or run 1.4.0

04 **Split drag-and-drop rearrangement on GTK** NEW 70

Splits can now be dragged to any side of any other split, across tabs and windows, using a drag handle at the top of each split; an overlay shows where the split will land before you release.

Rearrange terminal splits on GTK by dragging — when a tab has two or more splits, grab the drag handle at the top of any split and drop it onto another split, including in other tabs or windows.

> 📍 1. Open a GTK Ghostty window with two or more splits. 2. Grab the drag handle that appears at the top of any split. 3. Drag it onto another split (in the same tab, a different tab, or a different window). 4. An overlay shows which side of the target split it will land on; release to drop.

— Step-by-step UI path included 1.4.0

05 **Full Kitty graphics protocol coverage** NEW 60

Adds support for Kitty graphics animation and relative placements, achieving 100% coverage of the published Kitty graphics specification — making Ghostty the first non-Kitty terminal to reach full compliance.

— Names the spec milestone but no usage steps 1.4.0

THINNER COVERAGE BELOW

06 **macOS 27 support** NEW 30

Adds macOS 27 (Golden Gate) support while maintaining compatibility back to macOS 13 (Ventura).

— Just a compatibility note, no detail on changes 1.4.0

WAS THIS USEFUL? 👍 👎

Microsoft code-oss

Feedback ⓘ announced · Oct 26, 2026

BLOG ↗ RANK

The code-oss project is an open-source code editor combining editing, debugging, and extensibility for developers across Windows, macOS, and Linux.

> code-oss added configurable eagerness for GitHub Copilot inline suggestions, letting developers control how proactively suggestions appear and how large proposed edits are.

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

01 **Configurable eagerness for Copilot inline suggestions** NEW

45

GitHub Copilot inline suggestions now support three eagerness modes — low, medium, and high — which control how proactively Copilot offers suggestions and the size of the proposed edits.

.
.
.
.

— Names three modes but no setting key or UI path given. launch-20261026-5d40d6c6

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Oct 28, 2026 CHANGELOG ↗

RANK

Cloudflare is a connectivity cloud platform delivering internet security, performance, and developer infrastructure services.

> Cloudflare shipped jurisdiction pinning for Workflows, new OpenTelemetry-aligned span methods for Workers, SQL API introspection improvements, agent-workspace interrupt/resume APIs in `@cloudflare/computer`, a breaking migration for DNS nameserver type representation, and custom HTML/Liquid templates for Access pages.

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

01 Jurisdiction pinning for Workflows NEW

90

The `jurisdiction` field in `workflows` bindings in `wrangler.toml` / wrangler JSON config supports values `eu`, `us`, and `fedramp`, so Workflow instances only run and store data within that jurisdiction (requires Wrangler 4.152.0 or above). Jurisdiction is set at creation time and cannot be changed afterward; if an existing Workflow has a different jurisdiction, `wrangler deploy` and `wrangler versions upload` will stop before uploading.

Pin a Workflow to EU data centers to satisfy GDPR data-localization requirements.

toml

```toml
[[workflows]]
name = "my-workflow"
binding = "MY_WORKFLOW"
class_name = "MyWorkflow"
jurisdiction = "eu"
```

— Config key, allowed values, version requirement and breaking constraint all named

changelog-20261028-c0e5c7b3

02 include_attributes parameter for SQL API introspection NEW

90

The SQL API introspection endpoint (`GET https://api.cloudflare.com/client/v4/analytics/sql/introspection`) now accepts an `include_attributes` query parameter that returns discovered attribute keys and their types from the preceding seven days of data in an `attributes` response field.

Discover all attribute keys available in your HTTP requests dataset for use in SQL queries like `attributes['http.method']`.

```
$ curl --get
'https://api.cloudflare.com/client/v4/analytics/sql/introspection
' -H 'Authorization: Bearer <API_TOKEN>' --data-urlencode
'account_tag=<ACCOUNT_TAG>' --data-urlencode
'dataset_name=events.httpRequests' --data-urlencode
'include_attributes=true'
```

— Named endpoint and param with a working curl example `changelog-20261028-c0e5c7b3`

---

03 **Interrupt and resume support in agent workspaces** NEW

85

`@cloudflare/computer` (v0.7.0+) adds `onInterrupt: 'replay'` to `WorkerJavaScriptBackend`, allowing agent workspaces to resume interrupted runs after a deploy or crash by replaying recorded call results. Custom modules on `WorkerJavaScriptBackend` can call ctx.interrupt() to pause a run until an external approval decision is supplied via workspace.runtime.respond(), and new workspace.runtime.pending() and workspace.runtime.respond() APIs fetch and resolve pending approval calls.

— Names every API surface but lacks a runnable usage example `changelog-20261028-c0e5c7b3`

---

04 **Advanced Nameservers represented via nameservers.type**

85

BREAKING

DNS settings endpoints (`/zones/{zone_id}/dns_settings` and `/accounts/{account_id}/dns_settings`) add `nameservers.type: 'cloudflare.advanced'` to represent Advanced Nameservers, replacing the previous `foundation_dns: true` boolean. The `foundation_dns` boolean will stop being accepted in a future API change, requests mixing conflicting `nameservers.type` and `foundation_dns` values are currently rejected, and clients treating `nameservers.type` as a closed enum must be updated to recognize the new value before a seven-day rollout reaches their account.

— Named endpoints and migration deadline, but no code example `changelog-20261028-c0e5c7b3`

---

05 **OpenTelemetry-aligned span methods in Workers** NEW

75

Custom spans in Workers gain span.spanContext(), span.setStatus(status), and span.updateName(name) methods, aligning with the OpenTelemetry span API and enabling trace/span ID correlation, status export, and dynamic span renaming.

Return the active trace ID in an error response so clients and logs can be correlated to the same Workers trace.

```javascript
import { tracing } from 'cloudflare:workers';
export default async (request, env) => {
  const root = tracing.getActiveSpan();
  return tracing.startActiveSpan('handle', async (span) => {
    try {
      const response = await routeHandler(request, env);
      span.setStatus({ code: 'ok' });
      return response;
    } catch (err) {
      span.recordException(err);
      span.setStatus({ code: 'error', message: 'Route handler
failed' });
      return Response.json(
        { error: 'Internal error', traceId:
root?.spanContext().traceId },
        { status: 500 }
      );
    } finally {
      span.end();
    }
  });
};
```

— Named methods with a working code example for trace correlation

changelog-20261028-c0e5c7b3

THINNER COVERAGE BELOW

06 Custom HTML/Liquid templates for Access pages NEW

55

Cloudflare Access supports fully custom HTML and Liquid templates for user-facing pages (login, identity-denied, forbidden, and command-line approval), assignable per application or as account defaults from the Cloudflare One dashboard.

— Names page types and dashboard location, no code surface changelog-20261028-c0e5c7b3

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! A Workflow's `jurisdiction` is set at creation time and cannot be changed afterwards; if an existing Workflow has a different jurisdiction, `wrangler deploy` and `wrangler versions`

`upload` will stop before uploading.

! DNS settings endpoints will stop accepting the `foundation_dns` boolean in a future API change; clients must migrate writes to `nameservers.type: 'cloudflare.advanced'`. Currently, requests that include both `nameservers.type` and `foundation_dns` with conflicting values are rejected.

! Clients that treat `nameservers.type` as a closed enum must be updated to recognize the new `'cloudflare.advanced'` value or they will break when the seven-day rollout reaches their account.

WAS THIS USEFUL?  

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

ArmorCode

Feedback ⓘ announced · Oct 28, 2026 BLOG ↗ RANK

ArmorCode is the independent control plane where AI agents help find, prioritize, and fix vulnerabilities across apps, cloud, infra, and AI.

> ArmorCode launches Forge, a family of purpose-built LLMs — Assess, Verify, Decide, and Resolve — built on NVIDIA Nemotron and running inside Anya's Model Hub to power targeted exposure management decisions, alongside a new design partner program for customers to help shape future models.

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

01 **ArmorCode Forge LLM family inside Anya** NEW 60

ArmorCode Forge is a family of purpose-built LLMs built on NVIDIA Nemotron, running inside Anya's Model Hub alongside frontier models for targeted exposure management decisions. Each Forge decision includes a confidence level and reasoning; early versions add tags and comments for review rather than automatically changing a finding's status or severity.

— Names model foundation and hosting but no exact command or API.

launch-20261028-abfb2359

THINNER COVERAGE BELOW

02 **Forge Verify false positive detection** NEW 53

Forge Verify flags false positives by evaluating component applicability, runtime conditions, and scanner signals, returning an undetermined result when evidence is insufficient.

— Clear evaluation criteria but no concrete interface detail given.

launch-20261028-abfb2359

03 **Forge Resolve fix recommendation and validation** NEW 49

Forge Resolve recommends fixes, offers mitigations when a fix must wait, and explains how to validate the result, accounting for existing code, ownership, and controls.

— Describes scope well but lacks a concrete usage pointer. launch-20261028-abfb2359

04 **Forge Assess contextual severity recommendation** NEW 46

Forge Assess recommends contextual severity using asset, exposure, and business context while preserving the original scanner score.

— Describes mechanism but no UI path or command to invoke it. launch-20261028-abfb2359

05 **Forge Decide true-positive confirmation** NEW 43

Forge Decide evaluates available evidence to confirm true positives and produces a shareable rationale for remediation owners.

— Describes output but not how to access or trigger it. launch-20261028-abfb2359

06 **Forge design partner program** NEW 30

ArmorCode Forge design partner program allows customers to opt in to contributing training data and shaping future Forge models.

— Names the program but gives no enrollment mechanism or details.

launch-20261028-abfb2359

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 28, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue