THE CYBER TOOLCHAIN

NO. 132

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 6, 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 144 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 41 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 | Opus 5

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 today: Huntress can now push Microsoft 365 Conditional Access blocks for confirmed identity-attack

vim ? help

infrastructure, Sysdig Secure puts policy and session recording around live AI agents, and cve-lite-cli traces transitive findings back to parent upgrades.

RESPOND

Block confirmed identity campaigns across Microsoft 365 tenants without hand-building policies

Huntress now automatically deploys Conditional Access blocking policies against confirmed identity-attack campaigns in Managed ITDR. That turns campaign confirmation into tenant-level containment work instead of a manual policy authoring and rollout task.

Huntress

GOVERN

Put enforceable guardrails and audit trails around AI agents

Sysdig Secure launched AI Defense with live inventory, policy enforcement, session recording, spend attribution, and OWASP-aligned templates for AI agents. Teams get a control plane for agent activity instead of relying on ad hoc logs and after-the-fact reviews.

Sysdig Secure

PROTECT

Trace vulnerable transitive packages to the upgrade that actually removes them

cve-lite-cli now reports parent-upgrade chains for transitive findings and separates dev from production dependencies in SBOMs. That helps teams avoid chasing a buried package directly when the practical fix is upgrading the dependency that pulled it in.

cve-lite-cli

PROTECT

Reduce false vulnerable findings for commercial Spring fixes

Snyk Open Source taught its pURL API to recognize Broadcom Tanzu Spring commercial fixes. Enterprise Spring users can stop treating a component as exposed when their commercial distribution has already supplied the fix.

Snyk Open Source

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

web-check

Feedback ⓘ 2.4.0 NOTES ↗

All-in-one OSINT tool for analysing any website

web-check adds an expanded check view with copy buttons, light mode, and accessibility settings.

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/lissy93/web-check:2.4.0
```

git

```
$ git clone --branch 2.4.0 https://github.com/lissy93/web-check.git
# already have the repo? check out this version:
$ git checkout 2.4.0
```

published under the project's GitHub account · checked 6 Nov Wrong command?

web-check adds an expanded check view and copy buttons to results.

web-check adds light mode and accessibility settings.

NEW DETECTION CAA Records and mixed content checks

NEW INTEGRATION NetAPI compromised feed in threats check

UPDATED DETECTION Default DKIM selectors expanded

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Bugcrowd

Feedback ⓘ

Periodic Submissions Report: Triage Metrics Split by Priority (2026-11-05)

BLOG ↗ RANK

Bugcrowd is a crowdsourced security platform that connects organizations with ethical hackers to find and report vulnerabilities through bug bounty and penetration testing programs.

> Bugcrowd restructured its Periodic Submissions Report to break triage metrics out by priority, while dropping previous-month figures from the same view.

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

01 Priority-split triage metrics in Periodic Submissions Report

60

BREAKING

The Periodic Submissions Report now shows current-month submission counts and 'average days to triage for the current month' split into two groups — P1, and P2–P5 & NULL — so P1 triage performance is no longer averaged together with lower-priority work. As part of this change, the 'Avg. number of days to triage' fields for the previous month and current month were removed and previous-month triage metrics are no longer included in the report; use the date filter to view prior-month data instead.

Review P1-only triage performance for the current month without it being diluted by lower-priority averages.

> 📍 In the Periodic Submissions Report, locate the new P1 section showing 'submissions for the current month' and 'average days to triage for the current month', then compare against the P2-P5 & NULL section to benchmark priority-level triage speed.

— Names exact report fields and priority groups but no API/export detail

Periodic Submissions Report: Triage Met…

└──▷ **BREAKING ON UPGRADE**

! The 'Avg. number of days to triage' fields for the previous month and current month have been removed from the Periodic Submissions Report and replaced with priority-grouped current-month fields; previous-month triage metrics are no longer included.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

GitHub Feedback ⓘ

Usage metrics API adds code review details and addressed rate (2026-11-06)

BLOG ↗ RANK

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

> GitHub expanded the Copilot usage metrics API with granular code-review and telemetry/version-compliance reporting, and shipped a public-preview pull request Overview page that consolidates merge readiness and next actions.

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

01 **Code review metrics in Copilot usage API** NEW 95

Enterprise and organization `repos-1-day` reports gain a `copilot_code_review` object with `total_comments_evaluated` and `total_comments_addressed` fields to compute how often developers act on Copilot review feedback per repository. Comment counts can be broken down via `review_comments_by_severity` (`high`, `medium`, `low`), `review_comments_by_review_tier` (`lite`, `balanced`, `max`), `review_comments_by_comment_type` (`bug`, `security`, `custom`), and `review_comments_by_customization`, and a new `total_merged_no_human_change_request_after_review` field measures PRs merged without a subsequent human change request after Copilot review.

Compute the per-repository Copilot code review addressed rate across your organization to see how often developers act on feedback.

```
$ curl -H 'Authorization: Bearer <token>' -H 'Accept:
application/vnd.github+json'
'https://api.github.com/orgs/<org>/copilot/metrics/repos?
report=repos-1-day' | jq '.[] | {repo: .repository,
addressed_rate: (.copilot_code_review.total_comments_addressed /
.copilot_code_review.total_comments_evaluated)}'
```

Filter repos where Copilot reviews rarely prompt a human follow-up change request, indicating high reviewer confidence in Copilot's output.

```
$ curl -H 'Authorization: Bearer <token>' -H 'Accept:
application/vnd.github+json'
'https://api.github.com/enterprises/<enterprise>/copilot/metrics/
repos?report=repos-1-day' | jq '[.[] |
select(.copilot_code_review.total_merged_no_human_change_request_
after_review > 0) | {repo: .repository, clean_merges:
.copilot_code_review.total_merged_no_human_change_request_after_r
eview}]'
```

— Names every field and gives runnable curl/jq examples

Usage metrics API adds code review deta…

02 Telemetry and version-compliance fields in usage API NEW

93

Enterprise and organization user reports (1-day and 28-day) add a `client_telemetry_present` boolean to find users missing editor telemetry, and a `usage_metrics_requirements` field listing each IDE and Copilot extension as `met`, `unmet`, or `unknown` against minimum supported versions. New `missing_client_telemetry_user_count` and `below_usage_metrics_minimum_user_count` fields give daily counts in enterprise/organization 1-day reports and in each `day_totals` entry of 28-day reports.

Check the enterprise-level daily counts to gauge the blast radius of telemetry and version gaps across your whole org fleet before triaging individual users.

```
$ curl -H 'Authorization: Bearer <token>' -H 'Accept:
application/vnd.github+json'
'https://api.github.com/enterprises/<enterprise>/copilot/metrics/
usage?date=2026-11-05' | jq
'{missing_client_telemetry_user_count,
below_usage_metrics_minimum_user_count}'
```

— Specific field names with a runnable API example

Usage metrics API flags telemetry and v…

03 New pull request Overview page (public preview) NEW

77

A compact readiness summary now sits at the top of pull requests showing merge status, checks, approvals, and available next actions, replacing the need to scroll to the classic merge box; from the header, users can mark a draft PR ready for review or close it, depending on permissions. The page adds conversation filters for unresolved threads or comments you wrote, newest-to-oldest comment sorting, and an upper-right activity

summary for review comments, review requests, and recent updates. Opt in via the `Try the Overview page` button in the PR header, with an option to return to the classic view from the same menu.

Opt in to the new pull request Overview page during public preview to see merge readiness and next actions at the top.

> 📍 In a pull request, select the ✨ `Try the Overview page` button in the header. To return to the classic view, use the same menu.

Narrow a large pull request conversation to only unresolved threads so you can focus on what still needs attention.

> 📍 On the Overview page, use the conversation filter to select `Unresolved conversations`, then follow up on each open thread.

— Clear UI opt-in path given, but no API or config surface

`New pull request overview page in publi…`

NEW DETECTION SpaceX Starlink Public API v2 Secret detector

NEW INTEGRATION SpaceX secret scanning partner

UPDATED DETECTION Reduced false positives for Alibaba, AWS, Google, OpenVSX, Pydantic Services, and Tencent detectors

> `Secret scanning upgrades detectors, adds SpaceX as partner (2026-11-05)`
>
> **BLOG ↗** **STABLE**

WAS THIS USEFUL? 👍 👎

cve-lite-cli

Feedback ⓘ v1.41.0 NOTES ↗ RANK

Fast, developer-friendly JS/TS dependency vulnerability scanner with local lockfile scanning, OSV matching, direct vs transitive visibility, --fix, JSON output, and practical remediation guidance.

> cve-lite-cli v1.41.0 adds stdout and stable-filename report output, a redacted diagnostic bundle, parent-upgrade chains for transitive findings, and dev/production distinction in SBOMs, alongside two breaking changes around Node.js version and SARIF paths.

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

01 **Debug diagnostic bundle via --debug** NEW 80

`--debug` collects a diagnostic bundle for issue reports, producing an environment record, a machine-readable event stream, and a project-shape summary, with home paths, tokens, and private registry hosts redacted.

Capture a full diagnostic bundle to attach to a bug report when a scan produces unexpected results.

```
$ cve-lite /path/to/project --debug
```

— Explains bundle contents and redaction with a runnable command. v1.41.0

02 **Dev vs production dependency scope in SBOMs** NEW 80

SBOMs now distinguish development from production dependencies: SPDX output marks dev-only edges as `DEV_DEPENDENCY_OF`, and CycloneDX components carry a `scope` field (`optional` for dev, `required` for production).

— Names exact SPDX and CycloneDX fields and values. v1.41.0

03 **Stdout report output via --output -** NEW 70

Writing any report format to standard output is now possible by passing `--output -`, which keeps banners and messages off the stream; requesting more than one format at once with `--output -` is rejected.

Pipe a SARIF report directly into another tool in CI without writing a temporary file.

```
$ cve-lite /path/to/project --sarif --output -
```

— Names exact flag and example command, but behavior is brief. v1.41.0

04 **Stable filename output for CI** NEW 70

Passing a file path (not just a directory) to `--output` writes the report to a stable, predictable filename for CI; a path ending in a recognised report extension, or one that already exists as a file, is written to directly.

Write a JSON report to a predictable filename so a downstream CI step can always find it at the same path.

```
$ cve-lite /path/to/project --json --output reports/cve-report.json
```

— Flag and resolution logic named with a runnable example. v1.41.0

---

05 **SARIF output includes subdirectory folder path** BREAKING 60

SARIF output from a subdirectory scan now includes the folder path in locations; any post-processing step that adds that prefix manually will produce doubled paths and must be removed.

— Explains impact and required remediation step clearly. v1.41.0

THINNER COVERAGE BELOW

06 **Parent-upgrade chain for transitive findings** IMPROVED 55

Scan results now carry the verified parent-upgrade chain for a transitive finding, naming which direct dependency to move, from which version to which, and the resulting safe version.

— Describes mechanism but no flag or command shown. v1.41.0

---

07 **Minimum Node.js version raised to 22** BREAKING 50

The minimum supported Node.js version moves from 20 to 22; installations running Node 20 will no longer work.

— Clear before/after but no migration steps given. v1.41.0

└──▷ BREAKING ON UPGRADE

! The minimum supported Node.js version moves from 20 to 22; installations running Node 20 will no longer work.

! SARIF output from a subdirectory scan now includes the folder path in locations; any post-processing step that adds that prefix manually will produce doubled paths and must be removed.

WAS THIS USEFUL?  

Veracode

Feedback ⓘ CLI updates (2026-11-05) 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 CLI gained static-scan commands, async SCA remediation, RBAC-scoped container/IaC scans, and a wide set of new packaging flags, while the platform added v2 Reporting API filters and fields, a Custom PDF REST API, OAuth Client Credentials for XML APIs, and several breaking removals including the `/healthcheck/status` endpoint and legacy Azure DevOps and Jira integrations.

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

01 **Expanded packaging flags and detection in veracode package**

NEW 85

`veracode package` gains `--csharp-source` to package C# source without compiling, `--java-source` (or `VERACODE_PACKAGE_JAVA_SOURCE=true`) to package Java source, `--skip-packagers` to exclude artifacts such as unsupported SQL, `--exclude` to exclude specific Java packages or JAR files, `VERACODE_PACKAGE_GOLANG_GENERATE` to enable CGO code generation for Go, and `SRCLR_MAKE_WINDOWS` to preprocess C/C++ source with Make using a native Windows gcc or clang compiler. It also adds a generic packager for Apex, support for `requirements-dev.txt` Python development dependencies and the Poetry package manager, new exit codes 5, 6, and 7 for CI/CD automation, and improved toolchain detection in `veracode package discover` for JavaScript (npm/Yarn), Java, and Gradle projects.

– Names every flag and env var, but no example invocation shown CLI updates (2026-11-05)

02 **IaC scan policy gating in pipeline config** NEW 80

IaC scans can now be gated on a named policy via `iac: policy: 'My IaC Policy'` in `veracode.yml`, causing the pipeline job to fail when findings violate that policy.

Gate an IaC scan on a named policy in a GitLab pipeline so the job fails when findings violate that policy.

yaml

```yaml
# veracode.yml
iac:
  policy: 'My IaC Policy'
```

– Exact config key and working example provided Integration updates (2026-11-05)

03 New date filters on v2 findings reporting endpoint IMPROVED

80

The v2 findings endpoint `POST /analytics/findings/report` now accepts new date filters `first_found_start_date`, `first_found_end_date`, `resolved_start_date`, and `resolved_end_date` to narrow results within the still-required `last_updated_start_date` and `last_updated_end_date` range.

Request a findings report scoped to items first found in October, without client-side post-filtering, using the new v2 endpoint.

```
$ curl -X POST
'https://api.veracode.com/analytics/findings/report' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{
    "last_updated_start_date": "2026-10-01",
    "last_updated_end_date": "2026-10-31",
    "first_found_start_date": "2026-10-01",
    "first_found_end_date": "2026-10-31"
  }'
```

— Endpoint, filters, and runnable curl example given Platform updates (2026-11-05)

04 New fields and filters across the Reporting API NEW

65

The Reporting API gains several new fields and filters: `collections_list` in the Finding Reporting API response returns collection names associated with the application; `most_recent_scan` in the Scans report, when set to `true`, returns only the most recent scan per asset to reduce client-side filtering; `custom_severity` filters the Findings report; `data_refresh_date` indicates the most recent time Analytics data was refreshed; `library_instance_ref` uniquely identifies the library instance associated with a finding; and `action_type` filters the Audit Log report by specific user actions.

— Six named fields/filters, no example calls Platform updates (2026-11-05)

05 Azure DevOps Flaw Importer v3 deprecated DEPRECATED

60

Veracode Azure DevOps Extension v4.0.0 deprecates the Veracode Flaw Importer v3 task, replacing it with VeracodeFlawImporter v4; existing pipelines must be manually migrated from v3 to v4 for both Azure DevOps Server (upgrading from v3.35.0) and Azure DevOps Cloud.

— Names old/new task, lacks migration steps Integration updates (2026-11-05)

THINNER COVERAGE BELOW

06 **Static scan CLI commands** NEW 55

New `veracode scan static` command runs static scans and reviews security findings within a Git workflow; new `veracode findings static` command accesses cached static scan findings stored in the local Veracode directory.

— Named commands but no usage example given CLI updates (2026-11-05)

07 **Async Fire-and-Forget mode for SCA remediation** NEW 55

`veracode fix sca` gains an asynchronous Fire-and-Forget mode that submits SCA remediation requests for background processing without waiting for fixes or pull requests to complete.

— Named command and mode, no runnable example CLI updates (2026-11-05)

08 **Healthcheck endpoint removed** BREAKING 55

The `/healthcheck/status` endpoint has been permanently removed and now returns 404; use the `v2/principal` endpoint in the Identity REST API to check authenticated connectivity instead.

— Exact old and replacement endpoints named Platform updates (2026-11-05)

09 **OAuth Client Credentials for XML APIs** NEW 50

XML APIs now support OAuth Client Credentials authentication alongside the existing REST API support, and OAuth 2.0 Client credentials for Veracode integrations and APIs can be created and managed on the API Credentials page in the Veracode Platform.

— UI path named, no endpoint or flag Platform updates (2026-11-05)

10 **Target UUID triggering for dynamic scans** IMPROVED 45

The `veracode dynamic scan` command can now trigger dynamic scans by target UUID, in addition to the existing webhook UUID option.

— Command named, no example invocation CLI updates (2026-11-05)

11 **RBAC for container and IaC scans** NEW 40

New RBAC support restricts `veracode scan` container and IaC scans to application profiles associated with the requesting user's teams.

— Mechanism named but no flag or config shown CLI updates (2026-11-05)

12 **OAuth region switching in VS Code extension** IMPROVED 40

Veracode Scan for VS Code v1.16.3 adds an option to change the region during OAuth authentication, and displays the current or most recently used region below `Authentication` in the extension settings.

.
.
.
.

— UI location given, no exact setting name Integration updates (2026-11-05)

---

13 **Jira 9.x integration support ending** DEPRECATED 40

Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends on November 30, 2026; customers must upgrade to Jira 10 or later to continue using a supported version.

— Deadline given, no migration steps Integration updates (2026-11-05)

---

14 **report_type removed from v2 Reporting API requests** BREAKING 35

`report_type` is no longer accepted in the request body of the v2 Reporting API endpoints.

— Names removed field, no migration detail Platform updates (2026-11-05)

---

15 **GitHub Workflow autopackaging for Apex, VB6, Perl** IMPROVED 30

GitHub Workflow Integration v2.0.0 adds autopackaging support for Apex, VB6, and Perl.

— Names languages but no usage detail Integration updates (2026-11-05)

---

16 **Custom PDF REST API** NEW 25

New Custom PDF REST API lets you programmatically generate custom PDF reports.

— Only a name given, no endpoint or example Platform updates (2026-11-05)

---

17 **Firewall success events removed from Analytics** BREAKING 25

Success events have been removed from the Firewall dataset in Veracode Analytics.

— States removal only, no detail Platform updates (2026-11-05)

└──▷ **BREAKING ON UPGRADE**

! Veracode Azure DevOps Extension v4.0.0 deprecates the `Veracode Flaw Importer v3` task and replaces it with `VeracodeFlawImporter v4`; existing pipelines must be manually migrated from v3 to v4 for both Azure DevOps Server (upgrading from v3.35.0) and Azure DevOps Cloud.

! Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends on November 30, 2026; upgrade to Jira 10 or later to continue using a supported version.

! The `/healthcheck/status` endpoint now returns 404 and has been permanently removed; use the `v2/principal` endpoint in the Identity REST API to check authenticated connectivity.

! `report_type` is no longer accepted in the request body of the v2 Reporting API endpoints.

! Success events have been removed from the Firewall dataset in Veracode Analytics.

WAS THIS USEFUL? 👍 👎

Snyk Open Source

Feedback ⓘ changelog · Nov 4, 2026 BLOG ↗

RANK

Snyk Open Source is a developer security tool that scans project dependencies for known vulnerabilities and license issues, suggesting fixes or upgrades to remediate them.

> Snyk Open Source improved vulnerability accuracy for enterprise Spring users by teaching its pURL API to recognize Broadcom Tanzu Spring commercial fixes.

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

01 **Broadcom Tanzu Spring fix recognition in pURL API** IMPROVED

59

Snyk's pURL API (`GET` / `POST`) now recognizes Broadcom Tanzu Spring enterprise subscription builds as fixed, reducing false positives for mutual customers whose commercial Spring versions are already patched. No configuration changes are needed, but access must be enabled per organization by the Snyk account team.

— Names API and mechanism but access is manual/account-team gated 2026-11-04-b9a3c123

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

kingfisher

Feedback ⓘ v2.17.0 NOTES ↗ RANK

Detect secrets + live validation. Map blast radius.

> Kingfisher v2.17.0 extends credential-scoped revocation to eight more providers, ships a new JavaScript/TypeScript SDK, and enriches validation responses with identity metadata for several providers.

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

01 **Credential revocation expanded to eight providers** NEW

88

Kingfisher can now revoke credentials for Notion OAuth, Resend, Twilio key SIDs, Vercel AI Gateway, GCP API keys, Anthropic ordinary keys, OpenAI project/admin keys, and GitHub App user-token reports, using `--var` to supply management credentials, IDs and full-token mappings. Unsupported families and the distinction between accepted-versus-completed revocation results are documented in `crates/kingfisher-rules/data/revocation-sources.md`.

— Names flag, doc path and all eight providers v2.17.0

02 **JavaScript/TypeScript SDK preview** NEW

68

A native `kingfisher-secret-scanner` JavaScript/TypeScript SDK (0.1.0) is now available as an npm preview, with typed subpaths for rules, scanning, local inputs, Git, validation, analysis and explicit revocation. It supports CommonJS and ESM on Node 22.13+, 24 and 26.

— Names package, version and supported Node versions v2.17.0

03 **Identity metadata in validation responses** IMPROVED

62

Validation responses for AWS, Hugging Face, GitHub, GitLab, CircleCI, and Fastly now include allowlisted identity metadata without requiring additional requests; AWS responses specifically include the caller ARN, account ID, and user ID.

— Names providers and AWS-specific fields returned v2.17.0

THINNER COVERAGE BELOW

04 **Report viewer import format labeling** IMPROVED

25

The report viewer's upload instructions now name the supported Kingfisher, Gitleaks and TruffleHog report import formats.

— Minor UI text change, no mechanism given v2.17.0

WAS THIS USEFUL? 👍 👎

openbao

Feedback ⓘ CVE-2026-115356 CVE ↗

OpenBao is a software solution to manage, store, and distribute sensitive data including secrets, certificates, and keys.

`CVE-2026-115356`: OpenBao's Kubernetes Auth Method Allows Renewal on Revoked Tokens (CVSS 5.4, affects `<2.7.1, <2.6.4`)

WAS THIS USEFUL? 👍 👎

◆ **Software Supply Chain Security**

Anchore syft

Feedback ⓘ v1.57.0 NOTES ↗ RANK

CLI tool and library for generating a Software Bill of Materials from container images and filesystems

> Syft v1.57.0 adds GNU ELF build ID capture for cataloged binaries and introduces lazy lower-layer squashing to speed up targeted image scans.

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

01 **Lazy lower-layer squashing for scans** IMPROVED 35

Syft now builds lower-layer squashes only when a scan actually reads them, reducing unnecessary work for targeted image scans.

— Explains mechanism and benefit but no flag or metric v1.57.0

02 **GNU ELF build ID capture** NEW 32

Syft now captures GNU ELF build IDs when cataloging binaries, surfacing them in generated SBOMs.

— Names the data captured but no SBOM field or format shown v1.57.0

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ

Exempt specific packages from your cooldown policies (2026-11-05) BLOG ↗

RANK

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

> Cloudsmith's cooldown policies now support per-package exemptions, letting specific releases bypass the waiting period without disabling the policy for everything else.

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

01 **Package exemptions for cooldown policies** NEW 55

Specific packages can now be exempted from cooldown policies so they become immediately available while still being evaluated against all other active policies. In the Cloudsmith console, users open the cooldown policy and add an inline exemption for the package they want to release immediately.

Let a critical patch release bypass the cooldown period without disabling the policy for all other packages.

> 📍 In the Cloudsmith console, open your cooldown policy, then add an exemption for the specific package you want to release immediately.

— Clear UI steps given, but no API/flag or config key named.

Exempt specific packages from your cool…

WAS THIS USEFUL? 👍 👎

Socket

Feedback ⓘ New: Filter Alerts by Creation Date (2026-11-04)

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 added date-based filtering to both the Active Alerts page and Full Scans, letting teams narrow results to recent time windows and persist those views.

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

01 Date filtering for Active Alerts and Full Scans NEW

70

The Active Alerts page gains a Created date filter with preset ranges (Today, last 7, 14, 30, or 90 days) or a custom calendar range supporting between, since, and through modes, with filters preserved in URLs and saved views. The Full Scans date filter now supports the same presets (Today, last 7, 14, 30, or 90 days) and custom date range options.

Return to alerts created in the last 30 days by bookmarking the filtered URL — the Created date filter is preserved so the view reopens with the same range applied.

> 📍 In the Socket console, go to Active Alerts, open the Created date filter, select Last 30 days, then bookmark or save the resulting URL as a saved view.

— Names exact presets and modes; UI path given but no API/config key.

New: Filter Alerts by Creation Date (20…

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ Fastly Terraform Provider (beta) 0.2.8 (2026-11-05)

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 Terraform Provider (beta) gains new NGWAF resource query support and a safer legacy-to-auto migration workflow with preflight analysis.

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

01 **NGWAF Terraform query support** NEW 50

Fastly Terraform Provider now supports Terraform query operations (list resources and resource identities) for all NGWAF resources, including workspaces, account and workspace lists, alert integrations, rules, signals, redactions, thresholds, and virtual patches.

– Lists resource types but no example command shown

Fastly Terraform Provider (beta) 0.2.8 …

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Radare2

Feedback ⓘ 6.2.6 NOTES ↗ CODE ↗ RANK

UNIX-like reverse engineering framework and command-line toolset

> Radare2 6.2.6 adds global type propagation with a configurable trace limit, a new `afia` function-attribute command, richer syscall information in analysis JSON output, and new architecture-specific analysis support spanning V850, 16-bit ARM and ARM64.

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

01 Richer syscall and return-type detail in analysis JSON

IMPROVED 83

Improves analysis JSON output: `asj` now exposes syscall arity and argument formats, with Linux syscall clobbers and the MIPS error register recorded, while `afsj` distinguishes unknown return types from void so automated consumers can tell them apart.

Dump syscall argument layouts for every call site in a binary to feed into an audit script.

```
$ r2 -A /bin/target -c 'asj' | python3 -c 'import sys,json;
[print(e) for e in json.load(sys.stdin)]'
```

— Two named JSON commands with example pipeline 6.2.6

02 `afia` function attribute command and `is_inline` NEW

75

Adds the `afia` command to list and set function attributes, including a new `is_inline` attribute on `RAnalFunction`, so callers can be marked as inherited signatures for inlined helpers.

Inspect and tag an inlined helper so callers inherit the right signature after full analysis.

```
$ r2 -A /bin/target
> s sym.helper_fn
> afia          # list current attributes for the function
> afia is_inline 1  # mark the function as inline
```

— Named command and attribute with a runnable example 6.2.6

03 **Global type propagation and field inference** NEW 72

Adds global type propagation and field inference during analysis, including global discovery and typing from code accesses, plus a `types.maxtrace` config key to cap how many steps type propagation traces per function, preventing runaway analysis on complex binaries.

— Names config key but no usage example given 6.2.6

THINNER COVERAGE BELOW

04 **New calling conventions for V850 and ARM** NEW 58

Adds an IAR V850 calling convention with V850 stack-frame modeling, scalar argument and stack-variable recovery, and branch-predicate exposure for V850 binaries, plus the Swift calling convention to 16-bit ARM analysis.

— Mechanism described but no command or UI path Thin one-line addition, no usage detail 6.2.6

05 **HTTP referer and colon settings applied to POST** IMPROVED 58

Applies `http.referer` and `http.colon` settings to POST requests and matches the referer by prefix.

— Names exact config keys though no example given 6.2.6

06 **ARM64 checked-pointer arithmetic analysis** IMPROVED 35

Supports remaining ARM64 checked-pointer arithmetic and tagged-pointer subtraction in analysis.

— Names the behaviour but no mechanism or example 6.2.6

07 **Autoload API types via dependency index** IMPROVED 35

Autoloads API types through a library and OS dependency index.

— Describes mechanism loosely without named surface 6.2.6

08 **Interrupt handlers in function graph exports** IMPROVED 28

Function graph exports now include interrupt handlers.

— Bare statement with no mechanism or example 6.2.6

WAS THIS USEFUL?  

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-11-04 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds a `GET /apiv2/tasks/machine/<task_id>/` endpoint for distributed Guacamole VM resolution and disables expensive dashboard aggregations by default

└──▷ **USE IT**

Re-enable the top detections and top ASN dashboard widgets on a large deployment after opting back in.

custom/conf/web.conf

```
# custom/conf/web.conf
[statistics]
top_detections = yes
top_asn = yes
```

CAPEv2 adds a `GET /apiv2/tasks/machine/<task_id>/` API endpoint for distributed Guacamole VM resolution, with token-based authorization via `control_plane_token` from the `api` config section.

CAPEv2 now catches MongoDB aggregation timeouts (such as `pymongo.errors.ExecutionTimeout` / `MaxTimeMSExpired`) in `top_detections`, `top_asn`, and `get_stats_per_category`, logging a warning and caching an empty result so the dashboard renders instead of returning HTTP 500.

CAPEv2 now ships `systemd/00-ignore-vnet.network`, installed to `/etc/systemd/network/00-ignore-vnet.network`, to mark `vnet*`, `virbr*`, and `tap*` guest VM interfaces as `Unmanaged=yes` in `systemd-networkd`, preventing DHCP renewal timeouts on the host interface.

└──▷ **BREAKING ON UPGRADE**

! `top_detections` and `top_asn` are now disabled by default in `conf/default/web.conf.default`; existing deployments that relied on these dashboard widgets being enabled will need to explicitly re-enable them in their local `custom/conf/web.conf`.

WAS THIS USEFUL? 👍 👎

YARA-X

Feedback ⓘ v1.23.0 NOTES ↗ RANK

A rewrite of YARA in Rust.

> YARA-X v1.23.0 adds new regex word-boundary syntax, simplifies libmagic integration, and improves Unicode regex compilation performance.

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

01 **Half word-boundary regex assertions** NEW 75

YARA-X regular expressions now support `\b{start-half}` and `\b{end-half}` word boundary assertions, enabling half-boundary matching where a pattern only needs to align with the start or end edge of a word rather than both.

Match a pattern at the start of a word boundary half (e.g. the leading edge of a word) without requiring a full word boundary — useful for detecting tokens that begin but do not end a word.

yara

```
strings:
  $h = /\b{start-half}malware/
condition:
  $h
```

— Exact syntax and a runnable example given. v1.23.0

THINNER COVERAGE BELOW

02 **Auto-discovery of libmagic search path** IMPROVED 55

When building with the `magic-module` feature, YARA-X now auto-discovers the `libmagic` library search path, removing the need for manual path configuration.

— Names the build feature flag but no exact commands. v1.23.0

03 **UTF-8 compilation of Unicode character classes** IMPROVED 55

YARA-X now compiles Unicode character classes in regular expressions as UTF-8 byte sequences, producing smaller and faster PikeVM code.

— Explains mechanism but offers no user action. v1.23.0

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

Request Tracker / RTIR (Best Practical)

Feedback ⓘ

RT 6.0.4: Visualize and Navigate Lifecycles Right on the Ticket (2026-11-05)

BLOG ↗ RANK

Request tracker / rtir (best practical) is a open source ticketing system from Best Practical. Self-host RT or let us run it in the cloud.

> RT 6.0.4 adds an interactive lifecycle visualizer to ticket and asset pages and exposes a Notes field on lifecycle transitions via the REST API for AI agent instructions, alongside a Cytoscape-powered lifecycle editor.

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

01 Interactive lifecycle visualizer on tickets and assets NEW

72

Adds an interactive lifecycle visualizer component to ticket and asset pages, showing all statuses and transitions with click-to-change-status and hover descriptions. Lifecycle status descriptions and hover text are now configurable by RT administrators in the lifecycle configuration, so new users can learn the workflow directly from the ticket page. The component can be added via Admin > Queues > (select queue) > Page Layout, choosing its size and position on the page.



Add the lifecycle visualizer to a queue's ticket page so agents can see and change ticket status in one click.

📍 In the RT admin panel, go to Admin › Queues › (select queue) › Page Layout, then add the lifecycle visualizer component, choosing its size and position on the page.

— Names exact admin UI path and component behaviour

RT 6.0.4: Visualize and Navigate Lifecy…

THINNER COVERAGE BELOW

02 **Notes field on lifecycle transitions via REST API** NEW

45

An optional Notes field on lifecycle transitions is now exposed via the REST API, enabling additional instructions for AI agents that update tickets.

— Names API field but no endpoint or schema given

RT 6.0.4: Visualize and Navigate Lifecy…

03 **Cytoscape-based lifecycle editor** IMPROVED

35

The lifecycle editor now uses the Cytoscape library, giving administrators a smoother experience when building and editing lifecycles.

— Names library but no concrete behaviour change described

RT 6.0.4: Visualize and Navigate Lifecy…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ IDS / IPS

Cisco Snort3

Feedback ⓘ 3.13.2.0 NOTES ↗ RANK

Snort++

> Snort3 3.13.2.0 lets appid consume authoritative MCP inspector verdicts for application identification, adds control-channel file cache management, JSON stats output, and expanded http_inspect header enforcement alerting.

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

01 **Header validation and enforcement alerting in http_inspect**

IMPROVED 68

Snort3's `http_inspect` now alerts on header values exceeding maximum length, malformed headers, unsafe characters in header names and values, and disallowed request Content-Type values, and adds peg-counts for request enforcement events.

– Names all specific enforcement conditions and new peg-counts 3.13.2.0

02 **Control-channel file cache commands** NEW 64

Snort3's `file_api` adds control-channel commands to dump and clear the file cache via `file cache dump` control.

– Names concrete control-channel command usable immediately 3.13.2.0

THINNER COVERAGE BELOW

03 **MCP inspector verdicts in appid identification** NEW 48

Snort3's `appid` now consumes authoritative `MCP` inspector verdicts for application identification decisions and adds TLS host info logs.

– Names module and verdict source but no usage detail 3.13.2.0

04 **JSON format output for stats module** NEW 40

Snort3's `stats` module now supports JSON format output for dynamic stats.

– States new format but no config key or example given 3.13.2.0

05 **Per-flow memory tracking in flow module** NEW 38

Snort3's `flow` module adds per-flow memory tracking and reporting.

– Describes capability but not output format or access 3.13.2.0

06 **Trace logging for non-TCP stream inspectors** IMPROVED

38

Snort3's `stream` module now supports trace logging in non-TCP stream inspectors.

— Names module scope but not how to enable tracing 3.13.2.0

07 **Authentication and sandbox mode for control channel** NEW

38

Snort3's `control` channel adds authentication and sandbox mode.

— Names two new modes without configuration specifics 3.13.2.0

08 **PDF 2.0 processing in js_norm** NEW

30

Snort3's `js_norm` now enables PDF 2.0 processing.

— Bare statement with no mechanism or config detail 3.13.2.0

09 **Rate-filtered WEIRD logs in extractor** IMPROVED

30

Snort3's `extractor` now rate-filters WEIRD logs.

— Thin one-line addition with no parameters named 3.13.2.0

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Aviatrix CoPilot

Feedback ⓘ

New and Enhanced Features in 4.39.0 | NOTES ↗ | RANK

Aviatrix copilot is a 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 CoPilot 4.39.0 adds a preview custom rules engine and Terraform export for Access Control, namespace-based filtering for Kubernetes SmartGroups, and several smaller UI and connectivity improvements, alongside a breaking rename of External Groups.

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

01 External Groups renamed to ExternalIPGroups BREAKING

60

External Groups are renamed to `ExternalIPGroups` throughout CoPilot — any references to the old name in configurations, scripts, or integrations must be updated.

— Exact old/new names given, migration note included New and Enhanced Features in 4.39.0

THINNER COVERAGE BELOW

02 Access Control Enhancements NEW

55

CoPilot Access Control now supports custom rules (Preview) for fine-grained control over who can access CoPilot resources, and you can download Terraform code directly from Access Control rules to replicate access control policy as infrastructure-as-code.

— Names the capability but no rule syntax or scope given Clear UI action and IaC output named, no command shown

New and Enhanced Features in 4.39.0

03 BGP over LAN gateway details and reboot prompt IMPROVED

55

Gateway details now show HA LAN IP addresses for BGP over LAN connections, and a reboot prompt now appears when a BGP over LAN interface change requires a gateway reboot.

— Describes two concrete behavior changes with context

New and Enhanced Features in 4.39.0

04 Namespace label selector for Kubernetes SmartGroups NEW

50

Filter Kubernetes SmartGroups by namespace labels using the new namespace label selector in SmartGroup configuration.

– Names the exact config surface but no example values

New and Enhanced Features in 4.39.0

05 **SAML endpoint selection per controller** NEW 40

Select a SAML endpoint for the associated controller when configuring SAML authentication in CoPilot.

– Clear config action, but no detail on mechanism New and Enhanced Features in 4.39.0

06 **Retry failed gateway operations from UI** NEW 35

Retry failed gateway operations directly from the CoPilot UI.

– Clear UI action, no detail on which operations New and Enhanced Features in 4.39.0

07 **Recommended security control cards** NEW 20

Read-only recommended security control cards are now available in CoPilot.

– Bare description, no detail on content or location New and Enhanced Features in 4.39.0

08 **Telemetry export configuration options** IMPROVED 15

Telemetry export now has additional configuration options in the 4.39.0 release.

– No specifics on what options were added New and Enhanced Features in 4.39.0

└──▷ **ALSO FROM THESE RELEASES**

.
.
.

└──▷ **BREAKING ON UPGRADE**

! External Groups are renamed to `ExternalIPGroups` — any references to the old name in configurations, scripts, or integrations must be updated.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ announced · Nov 6, 2026 PRESS ↗

RANK

Sysdig Secure is a cloud security platform that provides runtime threat detection, vulnerability management, and compliance monitoring for containers, Kubernetes, and cloud workloads.

> Sysdig Secure launched AI Defense, a new runtime security and governance layer for AI agents that spans live inventory, policy enforcement, audit-grade session recording, spend attribution, and OWASP-aligned policy templates, now generally available on macOS and Windows.

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

01 **Live inventory and attribution of AI agents** NEW 63

AI Defense builds a live inventory of AI agents, harnesses, models, MCP servers, skills, and plugins from runtime telemetry, with every agent traceable to an accountable human across platforms including Claude Code, Codex, GitHub Copilot, and Cursor.

— Names platforms and scope but no UI path or command given. launch-20261106-55e985f3

02 **OWASP-aligned policy templates for LLM agents** NEW 60

AI Defense ships with policy templates including rules mapped to OWASP Top 10 for LLM Applications, with policies starting in alert mode and moving to enforcement when teams are ready.

— Names OWASP standard and alert-to-enforcement workflow as starting point.

launch-20261106-55e985f3

THINNER COVERAGE BELOW

03 **Runtime policy enforcement for AI agent actions** NEW 58

AI Defense enforces runtime policy against agent actions, blocking unvetted MCP servers, skills, and plugins before they run, and holding or blocking risky credential reads and behavior drift for human review, with exception paths for developers.

— Describes blocking mechanism but no concrete config or flag. launch-20261106-55e985f3

04 **Per-agent AI spend attribution** NEW 53

AI Defense attributes every model call to an agent, team, and individual — including calls that bypass an AI gateway — to surface real-time AI spend and identify where premium models are doing routine work.

— Explains attribution logic but no dashboard or query detail. launch-20261106-55e985f3

05 **Audit-grade agent session recording** NEW 48

AI Defense records every agent session as an immutable, audit-grade action graph showing each step in order, which policies were applied, when, and who signed off.

— Clear mechanism described but no access path noted. launch-20261106-55e985f3

06 **AI Defense general availability on macOS, Windows** NEW 25

Sysdig AI Defense is now generally available for macOS and Windows.

— Bare availability statement with no further detail. launch-20261106-55e985f3

WAS THIS USEFUL? 👍 👎

pike

Feedback ⓘ v1.0.14 NOTES ↗ RANK

Pike is a tool for determining the permissions or policy required for IAC code

> Pike's v1.0.14 release adds finer-grained permission splitting for destructive actions, credential-less CI baseline diffing, and expanded AWS/GCP policy analysis in the audit command.

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

01 **Credential-less CI diffing in compare** NEW 80

The `compare` subcommand gains `--write-baseline` to capture a permissions baseline without live cloud credentials, and `--target-file` to diff against that baseline in a later CI step.

Capture a permissions baseline in CI without live cloud credentials, then diff against it in a later pipeline step.

```
$ pike compare --write-baseline
pike compare --target-file baseline.json
```

— Two named flags with exact CI commands shown v1.0.14

02 **Sensitive permission splitting in scan** NEW 75

The new `--sensitive` flag on `pike scan` splits deletion-sensitive permissions into a separate policy, letting least-privilege reviews gate on destructive actions independently.

Separate high-risk deletion permissions into their own policy so least-privilege reviews can gate on destructive actions independently.

```
$ pike scan -d ./terraform --sensitive
```

— Names exact flag and subcommand with runnable example v1.0.14

THINNER COVERAGE BELOW

03 **AWS/GCP trust-policy and reference graph checks in audit**

IMPROVED

50

pike's `audit` command now checks AWS trust-policy (`assume_role_policy`) bindings, and generates attachment and policy-data reference graphs for AWS and GCP.

— Names surfaces checked but no example or flag given v1.0.14

WAS THIS USEFUL? 👍 👎

◆ **Container & Image Security**

kubescape

Feedback ⓘ v4.0.19 NOTES ↗ RANK

Kubescape is an open-source Kubernetes security platform for your IDE, CI/CD pipelines, and clusters. It includes risk analysis, security, compliance, and misconfiguration scanning, saving Kubernetes users and administrators precious time, effort, and resources.

> Kubescape's v4.0.19 release hardens the MCP server by restricting local IaC scans to a declared workspace boundary.

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

01 **Workspace boundary containment in MCP server** IMPROVED

45

Kubescape's `mcpserver` now enforces workspace boundary containment for local IaC scans, preventing the MCP server from accessing paths outside the declared workspace.

— Names the mechanism and component but no config detail or example v4.0.19

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ AWS Transfer Family servers in your inventory:

NOTES ↗ RANK

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

> Plerion expanded cloud inventory with AWS Transfer Family servers and broadened its end-of-life and secret-exposure detection checks across Azure Functions, GCP Compute Engine, GKE, and Cloud Functions.

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

01 **AWS Transfer Family server inventory** NEW 75

Plerion now collects AWS Transfer Family servers in inventory, showing endpoint type, protocols, identity provider, logging role, log destinations, security policy, and tags, with asset updates within minutes of a change when continuous scanning is on.

– Names all tracked fields but no direct command surface

AWS Transfer Family servers in your inv…

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Tencent Cloud CubeSandbox

Feedback ⓘ CVE-2026-120699

CVE ↗

Instant, Concurrent, Secure & Lightweight Sandbox for AI Agents.

`CVE-2026-120699`: Unrestricted host-directory bind-mount: any sandbox-create caller mounts arbitrary host paths (RW) into their MicroVM — host-root RCE, MicroVM escape, cross-tenant compromise (affects `<0.5.1`)

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Threat Intelligence

isMalicious

Feedback ⓘ updated · Nov 6, 2026 NOTES ↗ RANK

IsMalicious is a threat intelligence platform that checks IPs, domains, URLs, emails, phones and crypto wallets against indexed threat indicators.

> isMalicious added a dedicated email-scanning endpoint, HTTP Basic Auth support, mixed-indicator bulk checks, a newly-registered-domain feed, and dual-source certificate intelligence.

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

01 **Email scan endpoint for phishing and malware** NEW 85

A new `POST /mail/scan` endpoint scans inbound email for phishing and malware in a single call, returning a verdict, a recommended action, and reasons. It counts as one scan on the scan meter and never counts against the request quota.

Screen a raw inbound email for phishing or malware before it reaches a user's inbox, consuming one scan-meter credit.

```
$ curl -X POST https://api.ismalicious.com/mail/scan \
  -H 'X-API-KEY: <your-api-key>' \
  -H 'Content-Type: message/rfc822' \
  --data-binary @suspicious.eml
```

— Named endpoint, behaviour, quota mechanics, and a runnable example given.

snapshot-20261106

02 **Bulk IOC check endpoint** NEW 85

A new `/bulk/check` endpoint checks mixed batches of IPs, domains, and URLs in one request, accepting up to 500 indicators per request with responses streamed via chunked transfer encoding.

Submit a mixed batch of IPs, domains, and URLs for reputation checking in one round-trip, consuming up to 500 indicators.

```
$ curl -X POST https://api.ismalicious.com/bulk/check \
  -H 'Authorization: Basic <base64(apiKey:apiSecret)>' \
  -H 'Content-Type: application/json' \
  -d '{"indicators":["198.51.100.42","evil-
example.com","http://malware.example/payload"]}'
```

— Named endpoint, concrete limit, streaming behaviour, and runnable example.

snapshot-20261106

03 **Dual-source certificate facts** IMPROVED 70

isMalicious now reads `certificateFacts` from two independent sources — a live TLS handshake on port 443 for the certificate the site currently serves, and Certificate Transparency history from crt.sh — with unreachable sources reported as `unknown` rather than `none`.

— Field and both sources named, with explicit before/after reporting behaviour.

snapshot-20261106

04 **HTTP Basic Auth support** NEW 65

HTTP Basic Auth can now be used everywhere `X-API-KEY` is accepted, with `username=apiKey` and `password=apiSecret`. The TAXII/Sentinel form (any username, password = Base64 credential) is also supported.

— Exact credential fields named but no standalone example. snapshot-20261106

THINNER COVERAGE BELOW

05 **Newly-registered-domain feed** NEW 55

A daily NRD (newly registered domain) feed delivers domain age, registrar, nameserver, and isMalicious verdict, available in CSV and JSON.

— Fields and formats named, but no subscription command shown. snapshot-20261106

WAS THIS USEFUL?  

Feedly Threat Intelligence

Feedback ⓘ

`Campaign AI Models: Find new threat campaigns. Follow the ones that matter. (2026-11-05)`

BLOG ↗

Feedly threat intelligence is a collect, analyze, share actionable intel faster - Feedly TI

Feedly Threat Intelligence adds Campaign AI Models to discover and track named threat campaigns in AI Feeds, Intel Agents, Ask AI, Report Builder, and Automated Newsletters.

A `New Campaigns` AI Model surfaces articles about campaigns first detected in the past 30 days, combinable with an industry AI Model to focus discovery on a specific sector.

NEW DETECTION Individual campaign AI Models let analysts follow a named campaign across sectors for continuous collection of TTPs and IoCs, without waiting for threat-actor attribution.

Ask AI Research can reference a saved Org Profile item to connect a campaign's reported targets and access methods to your organization's environment.

Report Builder and Automated Newsletters accept campaign feed articles as source material, with AI Overview prompts supporting saved context items such as `@security leadership` and `@company overview`.

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ `CVE-2026-121708` CVE ↗

Open Cyber Threat Intelligence Platform

`CVE-2026-121708`: Broken Object Level Authorization (BOLA) for draft workspaces (CVSS 8.1, affects `< 7.260513.0`)

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ

Managed ITDR's New Capability, Adversary Infrastructure, Now in GA

NOTES ↗

RANK

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

> Huntress brought Adversary Infrastructure to general availability in Managed ITDR, automatically deploying Conditional Access blocking policies against confirmed identity-attack campaigns across Microsoft 365 tenants.

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

01 Adversary Infrastructure blocking in Managed ITDR NEW

83

Huntress now automatically writes and deploys one named location and one blocking Conditional Access policy per confirmed adversary campaign to every eligible, opted-in Microsoft 365 tenant the same day an attack is confirmed. Enable it via the Adversary Infrastructure toggle in ITDR settings, set a default at the account level, override per organization, or enable/disable in bulk; a new Adversary Infrastructure dashboard under ITDR in the left nav (available at account and organization level) shows covered tenants, campaign policies in force, and a table of every blocked sign-in. Every blocked sign-in generates a Low severity incident report naming the targeted identity, source network, and campaign. Tenants ineligible for Conditional Access (lacking Entra ID P1/P2, or with Security Defaults on) still get the same detections and incident reports, delivered after sign-in rather than before.

Opt an entire account into Adversary Infrastructure blocking so all eligible Microsoft 365 tenants automatically receive Conditional Access policies when new adversary campaigns are confirmed.

> 📍 In the Huntress portal, go to ITDR settings and turn on the Adversary Infrastructure toggle at the account level. To override for a specific organization, navigate to that organization's settings and adjust the toggle independently.

— Names toggle, dashboard path, policy mechanism and eligibility fallback.

Managed ITDR's New Capability, Adversar…

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Redesigned Help and Support panel (2026-11-05) 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 shipped a redesigned Help and Support panel with AI-generated documentation and consolidated support tools, plus improved visibility into in-progress external attack surface discovery.

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

01 Seed and root domains stay visible during asset discovery

IMPROVED 75

Seed and root domain records in Attack Surface Discovery now persist throughout the discovery lifecycle, remaining listed under Attack Surface Discovery → Internet-facing assets even when scan results are not yet available. Assets still collecting discovery data are marked with the Anchor stage icon, and WHOIS and DNS zone information is shown immediately on discovered assets before full enrichment completes, with open port, certificate, and scan details added as discovery finishes.

Review newly added seed and root domains before scan enrichment finishes, to get early WHOIS and DNS context on your external attack surface.

> 📍 In the console, go to Cyber Risk Exposure Management → Continuous Risk Management → Attack Surface Discovery → Internet-facing assets and look for entries marked with the Anchor stage icon to identify assets still collecting discovery data.

— Explains lifecycle mechanism and exact UI path, no API/flag.

Keep seed and root domains visible duri…

02 Redesigned Help and Support panel NEW

65

The Knowledge Hub tab now shows an AI-generated summary of the documentation for the current screen, with a Check Full Content option to open the full documentation in an in-console window whose search, language, and dark/light mode settings align to the console. A new Feedback & Support tab brings together feedback sharing, diagnostic testing, support case creation, and case history in one place, and a new More Resources tab links to training, the Automation Center, and service status.

Access AI-generated documentation for the screen you are currently on without leaving the console.

📍 1. Click the question mark icon in the console toolbar. 2. Select the `Knowledge Hub` tab — an AI-generated summary for the current screen appears automatically. 3. Click `Check Full Content` to open the full documentation in an in-console window with search and your language and dark/light mode settings applied.

Run a diagnostic test or create a support case without leaving the console.

📍 1. Click the question mark icon in the console toolbar. 2. Select the `Feedback & Support` tab. 3. Choose `Run a diagnostic test`, `Create a support case`, or `View case history` as needed.

— Names three tabs and their functions but no deeper mechanism.

`Redesigned Help and Support panel (2026…`

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Allow partner platforms in context-based ingress policies (2026-11-05) BLOG ↗

RANK

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

> Databricks expanded its Genie AI/BI surface this window with field display names, Unity Catalog relationship imports, access-scoped agent answers, a Monitor tab identity column, and OpenTelemetry trace export, while also adding Claude Code/Codex support inside SSH tunnels, a beta GitHub Action to publish Unity Gateway skills, and partner-platform allowances in context-based ingress policies.

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

01 **Display names for dashboard field variable options via Genie Code**

NEW 80

Each option in a field variable can now have a display name — for example `Total revenue` instead of `Sum of sales` — and Genie Code can set, rename, or clear these display names via natural-language commands while preserving ones already set.

When a field variable exposes raw column names, set readable labels so business users see `Total revenue` instead of `Sum of sales` without changing the underlying query.

> 📍 In the dashboard editor, open a field variable, select an option, and enter a display name in the `Display name` field. To bulk-update names, open Genie Code and type: 'Set the display name for the Revenue option to Total revenue'.

— Exact UI path plus a runnable Genie Code command example

Dashboard enhancements and bug fixes (2…

02 **Unity Catalog metric view relationship import into dashboards (Beta)**

NEW 75

Using Genie Code for dashboard authoring (Beta), relationships declared on Unity Catalog metric views can be imported into a dashboard as dashboard relationships, avoiding manual redefinition.

When authoring a dashboard over Unity Catalog metric views that already declare relationships, pull those relationships in automatically rather than redefining them by hand.

> 📍 In the dashboard editor, open Genie Code and type: 'Import relationships from Unity Catalog metric views'. Genie Code will add the declared relationships as dashboard relationships.

— Beta scope and mechanism named, with a runnable command example

Dashboard enhancements and bug fixes (2…

---

03 **Genie Agents access scoping and monitoring enhancements**

IMPROVED

75

Genie Agents now answer using only the volumes or MCP connections a user can access, so users with partial access are no longer blocked from using an agent. The Monitor tab adds a Custom claim column, enabled from the column picker, showing each message's OAuth custom identity claim to identify which end user sent it. Benchmark evaluations now run questions in parallel, finishing sooner.

Track which app end user asked each question when an OAuth custom identity claim is in use — open the Monitor tab and add the identity column.

> 📍 In the Genie Agent UI, go to the Monitor tab, open the column picker, and enable the Custom claim column to see each message's OAuth custom identity claim alongside the message.

— Names the column and access scopes, gives exact UI steps

Genie Agents enhancements (2026-11-05)

---

04 **Claude Code and Codex support in SSH tunnel sessions**

NEW INTEGRATION NEW

75

Claude Code and Codex coding agents can now run inside SSH tunnel sessions, with model requests routed through Unity Gateway and authenticated using the SSH session credentials. Requires Databricks CLI version `1.23.0` or above.

— Names the routing mechanism and exact CLI version requirement

Run Claude Code and Codex in the SSH tu…

---

05 **OpenTelemetry trace export for Genie Code and Genie One** NEW

70

Conversation turns from Genie Code and Genie One can be exported to the Unity Gateway unified trace table as OpenTelemetry spans. A workspace admin enables tracing per product from the Previews page, and by default only the trace table owner can query the traces.

— Mechanism and permissions detailed, admin UI path named but no exact command

Genie One enhancements (2026-11-05)

06 **GitHub Action to publish Unity Gateway Skills (Beta)** NEW

65

A new GitHub Action (beta) publishes Git-managed Unity Gateway skills to Unity Catalog automatically on merges to the publishing branch, by mapping folders to skills.

— Names the automation trigger and mapping mechanism, beta status noted

Publish Unity Gateway Skills with a Git…

THINNER COVERAGE BELOW

07 **Partner platforms allowed in context-based ingress policies**

NEW INTEGRATION NEW

55

Context-based ingress policies can now allow traffic from Power BI, Tableau Cloud, `dbt Cloud`, and the Genie apps for Slack and Microsoft Teams as recognized sources, without requiring manual IP address tracking.

— Names platforms but not the policy config surface itself

Allow partner platforms in context-base…

08 **Dashboard editor variable reordering and waterfall chart polish**

IMPROVED

35

Dashboard variables can now be reordered directly in the dashboard editor without workarounds, and waterfall chart labels remain visible when they are close to the top or bottom of the axis.

— Two thin UI fixes, no mechanism beyond the fix itself

Dashboard enhancements and bug fixes (2…

09 **Expanded visualization side panel in Genie One** NEW

35

Visualizations from a Genie One response can now be opened in a side panel to interact with them at a larger size next to the conversation.

— Describes the UI change but no further mechanism Genie One enhancements (2026-11-05)

└──▷ ALSO FROM THESE RELEASES

External knowledge sources are in Beta (2026-11-05) BLOG ↗ STABLE

Databricks adds external knowledge sources (Beta), enabling account admins to allow Genie to use Power BI measure definitions when answering questions.

> Enable Power BI semantic model metadata ingestion into Genie by turning on External knowledge from the account console Previews page, so Genie can use Power BI measure definitions when answering questions.

> Moonshot AI Kimi K3.1 Pro is now available on Unity Gateway (2026-11-04)
>
> BLOG ↗ STABLE

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Security Validation (BAS & Automated Pentest)

SafeBreach

Feedback ⓘ announced · Nov 6, 2026 PRESS ↗

SafeBreach is the leading AI CTEM platform for the enterprise—grounded in AEV to prioritize, prove, and reduce cyber risk

SafeBreach adds a Wiz Defend integration with an auditable attack-chain log that maps each simulated attack step to the Wiz Defend event it triggered.

NEW INTEGRATION SafeBreach now integrates with Wiz Defend via Wiz-validated APIs and permissions, correlating simulated attack activity against Wiz Defend cloud events and threat-detection issues matched on hostname and IP.

An auditable attack-chain log now records every step of a simulated attack and maps each action to the Wiz Defend event it produced, giving teams a traceable record of detection coverage.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta

Feedback ⓘ updated · Nov 6, 2026 NOTES ↗ RANK

Use Immuta to provision governed access to data.

> Immuta introduces Dynamic Review Flows, an AI-assisted decision tree builder for automating data access governance, including dedicated handling for agentic access requests.

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

01 **Dynamic Review Flows for access request automation** NEW

66

Dynamic Review Flows is a visual decision tree builder that branches on request type, consumer identity, request form answers, and AI-generated recommendations to automatically approve, deny, or route data access requests to a human reviewer. It also supports agentic access routing, evaluating consumer type and request context so agent requests meeting predefined criteria are auto-approved instantly while edge cases are escalated to a human reviewer. The feature introduces breaking API changes and follows a phased rollout: opt-in from 11/4–12/3, on by default 12/4–1/3, and enabled for all accounts from 1/4; early access requires contacting an Immuta representative or submitting a support ticket.

— Describes mechanism and rollout dates but no concrete config/API names given.

snapshot-20261106

└──▷ **BREAKING ON UPGRADE**

! Dynamic Review Flows introduces breaking API changes; the feature follows a phased rollout (opt-in 11/4–12/3, on by default 12/4–1/3, enabled for all accounts 1/4). Contact your Immuta representative or submit a support ticket to enable early.

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

SPIFFE spire

Feedback ⓘ v1.16.0 NOTES ↗ RANK

The SPIFFE Runtime Environment

> SPIRE v1.16.0 introduces experimental WIT-SVID issuance and a KMIP key manager plugin, adds Microsoft Entra ID authentication for Azure-hosted datastores, and ships a dozen-plus agent/server tuning options alongside several breaking changes to agent identity format and workload attestors.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 Experimental WIT-SVID issuance NEW

95

Workloads can fetch WIT-SVIDs via `spire-agent api fetch wit` once `enable_wit_svids` is set on both server and agent; TTLs are configurable per-entry or via the default `default_wit_svid_ttl`, a client-side cache is maintained, and WIT signing keys are published through the `spire` UpstreamAuthority plugin. Marked experimental.

Fetch a WIT-SVID for a workload once the experimental feature is enabled on both server and agent.

```
$ spire-agent api fetch wit
```

— Names CLI command, config keys, cache and signing mechanism v1.16.0

02 In-process log file rotation NEW

80

Setting `log_file_rotation` on the agent or server rotates log files without restarting the process, with support for external log file moves and deletes on Windows.

Enable in-process log rotation on the SPIRE agent so logs are cycled without restarting the process.

```hcl
agent {
  log_file          = "/var/log/spire/agent.log"
  log_file_rotation = true
}
```

— Exact config key plus runnable example v1.16.0

03 Node attestor selector and restriction options IMPROVED

80

Several node attestor plugins gained new selector or restriction options: `x509pop` adds `group_template` and `allowed_groups` for group selectors and `extension_oids` to expose arbitrary certificate extensions as selectors; `tpm_devid` can select workloads by TPM DevID subject attribute OID; `sshpop` and `http_challenge` add an optional `verify_client_ip` setting to restrict attestation to the calling client's IP address; and `azure_imds` adds `trust_bundle_path` to supply additional root CAs.

— Names every plugin and option, lacks usage example v1.16.0

---

04 **Agent-to-server connection tuning options** NEW 70

New experimental agent options tune agent-to-server behavior: `server_load_balancing_config` configures the gRPC load balancing policy for agent-to-server connections, and `sync_retry_backoff`, `rpc_timeout`, and `max_bundle_workers` tune agent sync behavior.

— Four named config keys but no example or defaults given v1.16.0

---

05 **Audit log exclusion for agent sync** NEW 70

Setting `audit_log_exclude_agent_sync` on the server omits successful bundle and authorized-entry read records, reducing audit log volume.

Suppress high-volume agent-sync audit records on the SPIRE server to keep audit logs focused on meaningful events.

hcl

```
server {
  audit_log_exclude_agent_sync = true
}
```

— Exact config key with runnable config example v1.16.0

---

06 **Agent SPIFFE ID path now enforced** BREAKING 70

Agent SPIFFE IDs must now reside under `/spire/agent`; attestations producing other paths fail unless the experimental option `allow_nonconforming_agent_ids` is set. This affects users of external plugins and custom `agent_path_template` values.

— Names path requirement, escape-hatch option and affected customizations v1.16.0

---

07 **Local-file federation bootstrap** NEW 65

Federation trust for `https_spiffe` can now be bootstrapped from a local file using the `bootstrap_bundle_path` option on a `federates_with` block.

— Names config block and option, no example shown v1.16.0

08 **Workload API as OIDC Discovery Provider cert source** NEW

65

The OIDC Discovery Provider can now serve HTTPS using an X509-SVID from the Workload API by configuring the new `serving_cert_source` section with its `workload_api` variant.

— Names config section and variant, no example v1.16.0

09 **Container locator changes for docker and k8s attestors**

65

BREAKING

Rootless Podman workload attestation now requires the `use_rootless_podman` option to be explicitly set in the `docker` workload attestor; existing rootless Podman setups will break without it. The mountinfo fallback in the container locator used by the `docker` and `k8s` workload attestors has been removed.

— Names both affected plugins and the removed/required mechanism v1.16.0

THINNER COVERAGE BELOW

10 **Exclude completed pods in k8s workload attestor** IMPROVED

55

Enabling `exclude_completed_pods` in the `k8s` workload attestor drops completed pods from the kubelet cache.

— Names config option but minimal mechanism detail v1.16.0

11 **ca_journals schema migration** BREAKING

55

The `ca_journals` table in the `sql` datastore gains a `journal_id` column; deployments with `disable_migration` set must plan for this schema migration.

— Names table, column and affected config flag v1.16.0

12 **KMIP KeyManager plugin** NEW

50

A new `kmip` KeyManager plugin lets SPIRE store its keys in KMIP-compliant external key servers.

— Names plugin but no config detail or example v1.16.0

13 **Microsoft Entra ID auth for Azure databases** NEW

50

The `sql` datastore can authenticate to Azure PostgreSQL and MySQL using Microsoft Entra ID instead of static credentials.

— Names databases and datastore but no config keys shown v1.16.0

14 **use_sync_authorized_entries option removed** BREAKING

45

The `use_sync_authorized_entries` agent option has been removed.

— Bare removal notice, no migration guidance given `v1.16.0`

---

15 **Broker API drops JWT-SVIDs for admin/downstream** BREAKING

45

The Broker API no longer serves JWT-SVIDs for admin or downstream entries.

— States the change but no migration path or reasoning `v1.16.0`

└──▷ **BREAKING ON UPGRADE**

! Agent SPIFFE IDs must now reside under `/spire/agent`; attestations producing other paths fail unless the experimental option `allow_nonconforming_agent_ids` is set — affects users of external plugins and custom `agent_path_template` values.

! Rootless Podman workload attestation now requires the `use_rootless_podman` option to be explicitly set in the `docker` workload attestor; existing rootless Podman setups will break without it.

! The mountinfo fallback in the container locator used by the `docker` and `k8s` workload attestors has been removed.

! The `use_sync_authorized_entries` agent option has been removed.

! The `ca_journals` table in the `sql` datastore gains a `journal_id` column; deployments with `disable_migration` set must plan for this schema migration.

! The Broker API no longer serves JWT-SVIDs for admin or downstream entries.

WAS THIS USEFUL?  

◆ Secure Access (VPN / SASE / ZTNA)

OpenVPN

Feedback ⓘ v2.7.9 NOTES ↗ CODE ↗ RANK

OpenVPN is an open source VPN daemon

> OpenVPN 2.7.9 improves username handling in the management interface password prompt and enables better Linux kernel packet prioritization on tun/tap devices.

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

01 Username hinting in management interface prompts IMPROVED

70

The management interface can now hint the username in `>PASSWORD '<UP_TYPE>'` prompts when a config has a username but no password set, letting GUIs pre-fill the username field. This is enabled when the management client announces compatibility level >= 6.

— Names exact prompt format and compatibility requirement. v2.7.9

02 IFF_BACKPRESSURE on Linux tun/tap devices IMPROVED

65

OpenVPN on Linux now sets `IFF_BACKPRESSURE` on tun/tap devices (when available), allowing the kernel to apply proper qdisc prioritization to incoming packets instead of dropping them when userspace is slow.

— Names the kernel flag and explains the mechanism and benefit. v2.7.9

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ

Public IP address device posture attribute (2026-11-05) BLOG ↗ RANK

Tailscale creates private WireGuard-based mesh networks with centralized identity and access controls.

> Tailscale graduated the public IP address device posture attribute to general availability, enabling access policies based on a device's public IP.

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

01 **Public IP posture attribute reaches GA** IMPROVED 70

The `ip:publicAddress` device posture attribute is now GA and can be used as a condition in posture policies; the value is also visible per-device on the Machines page of the admin console.

Restrict access to a resource based on a device's public IP address using the now-GA posture attribute.

> 📍 In the admin console, go to Machines, select a device to confirm its `ip:publicAddress` value, then open your posture policy and add a condition using `ip:publicAddress` to allow or deny access based on that address.

— Names exact attribute and UI path but lacks mechanism detail

Public IP address device posture attrib…

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Ping Identity

Feedback ⓘ 04 Nov 2026 (2026-11-05) BLOG ↗

RANK

Ping Identity is an identity and access management platform providing authentication, single sign-on, multi-factor authentication, and identity governance services.

> Ping Identity's Advanced Identity Cloud now lets organizations govern, provision, and reconcile multiple accounts per identity on non-authoritative provisioning applications, with new reconciliation and account-management UI, alongside a breaking change requiring migration of existing role assignments.

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

01 **Multiple accounts per identity in provisioning and reconciliation**

NEW 88

Advanced Identity Cloud can now govern, provision, and reconcile multiple accounts per identity on non-authoritative provisioning applications, letting organizations separate distinct user personas (for example an admin account versus a standard account) under independent lifecycles. Custom account types can be added in the `Account Types` section of the `Reconciliation > Settings` page, and setting `Found Already Linked` to `Onboard` on that same page allows multiple accounts to be linked during reconciliation instead of being rejected. Administrators can select a target account type when making targeted assignments — directly to users or through roles — to control which account receives the application. The `Accounts` tab now replaces the `Users` and `Roles` tabs on provisioning applications, adding `Change owner` and `Unlink` actions, and a new `Unlinked Accounts` tab under `Reconciliation` lets admins link orphan target accounts to identities or delete them.

Linking orphan target accounts discovered during reconciliation to the correct identity, rather than deleting them.

📍 In Advanced Identity Cloud, go to Reconciliation › Unlinked Accounts. Select an orphan account and choose `Link` to associate it with an existing identity, or `Delete` to remove it.

Separating a privileged admin persona from a standard user account for the same identity during reconciliation.

📍 In Advanced Identity Cloud, go to Reconciliation › Settings › Account Types and add a custom type (for example, `Admin`). Then set `Found Already Linked` to

Onboard so reconciliation links additional accounts. When assigning the application to a user or role, select the Admin account type as the target.

— Names every settings page and tab involved, with step-by-step examples.

04 Nov 2026 (2026-11-05)

---

02 **Migration required for existing role assignments** BREAKING

65

Pre-existing role assignments must be migrated to Default from the role's Applications tab before they apply to new accounts created under the multiple-account model. Account types cannot be changed once created; existing accounts retain their assigned type even if that type is later removed from the configuration.

— Names the exact UI tab for migration but no API or CLI step. 04 Nov 2026 (2026-11-05)

└──▷ **BREAKING ON UPGRADE**

! Pre-existing role assignments must be migrated to Default from the role's Applications tab before they apply to new accounts created under the multiple-account model.

! Account types cannot be changed once created; existing accounts retain their assigned type even if that type is removed from the configuration.

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ Ory Hydra, Ory Elements v26.4.12 released NOTES ↗

RANK

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

> Ory Hydra tightens control over token lifespan configuration for dynamically registered OpenID Connect clients, moving that setting exclusively to administrators.

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

01 Admin-only token lifespans for dynamic client registration

55

BREAKING

Ory Hydra now restricts token lifespan settings to administrators for dynamically registered OpenID Connect clients, removing per-client control via Dynamic Client Registration. Dynamically registered clients that previously set their own token lifespans will now inherit administrator-defined values instead.

— Clear before/after behavior but no exact config key or endpoint named.

Ory Hydra, Ory Elements v26.4.12 releas…

└──▷ **BREAKING ON UPGRADE**

! OpenID Connect Dynamic Client Registration no longer accepts per-client token lifespan settings — dynamically registered clients that previously set their own token lifespans will now inherit administrator-defined values.

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Version 2026.45.1 (2026-11-04) BLOG ↗ RANK

Okta identity engine is a secure, scalable, and highly available authentication and user management for any app.

> Okta Identity Engine's 2026.45.1 release adds support for staged AI agent import and the ability to configure multiple `AGENT_IMPORT` client authentication settings per app instance.

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

01 **Staged AI agent import and multi-auth settings** NEW

45

Okta Identity Engine can now import staged AI agents, and app instances can be configured with multiple `AGENT_IMPORT` client authentication settings rather than just one.

— Names config surface but lacks mechanism or UI path detail

Version 2026.45.1 (2026-11-04)

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 4th, 2026 NOTES ↗

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

Compyl's Risks module now loads faster for large risk registers, and the V2 API now returns clear errors for invalid role or nonexistent user assignments.

The Risks module now loads faster for large risk registers, with quicker risk opening, chart viewing, and navigation.

The V2 API now returns a clear error when creating or updating IT Assets, Contracts, Clients, or Incidents with an invalid role or nonexistent user.

UPDATED INTEGRATION Supabase integration — generally available

WAS THIS USEFUL? 👍 👎

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ RegScale 6.36.0.1 (2026-11-03) BLOG ↗ 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's only update this window was a scoped filter setting for the enterprise Risk Rollup status board, ensuring consistent filtering across all board views.

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

01 **Scoped filtering for Risk Rollup status board** IMPROVED

45

The Risk Rollup status board now supports a single scope setting that applies filters consistently across every view in the board, letting risk leads evaluate where risk is concentrated and which treatments are in play.

— Describes behavior and purpose but no concrete config name or steps.

RegScale 6.36.0.1 (2026-11-03)

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 3.2.0 (2026-11-06) BLOG ↗ RANK

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

> Runlayer's window centers on suggestion policies that nudge agents toward preferred tools, spend/run limit controls reachable via MCP clients, direct in-place artifact editing, expanded tunnel connector observability, and a breaking increase in Kubernetes operator memory requirements.

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

01 Custom agent/session identifiers for Claude Code OTLP sessions

NEW 90

Claude Code sessions sent over OpenTelemetry now carry custom agent identifiers: set `sidekick.agent_id` and `sidekick.session_key` via `OTEL_RESOURCE_ATTRIBUTES`, and they appear as External metadata on the session.

Attach your own agent and session identifiers to Claude Code OTLP sessions so you can trace a Runlayer session back to your internal agent.

```
$ export OTEL_RESOURCE_ATTRIBUTES="sidekick.agent_id=my-agent-42,sidekick.session_key=proj-session-7"
```

— Exact env var and attribute names given; fully actionable. 3.2.0 (2026-11-06)

02 Increased memory requirement for Kubernetes operator workers

90

BREAKING

The Kubernetes operator now requires at least 4096 MiB of memory per slow-pool worker on `components.topicCpu` (up from 2048 MiB); any `RunlayerInstance` with `topicCpu.gunicornWorkers` set to 3 or more without a matching `resources` memory override fails validation after upgrade and is not reconciled — including version bumps — until memory is raised or worker count is lowered.

— Exact config keys, values, and remediation path given. 3.1.0 (2026-11-05)

03 Suggestion policies to steer agent tool use NEW 80

Admins can create suggestion policies that steer agents toward preferred tools — for example, suggesting `create_app` instead of a third-party deploy — without blocking

the original tool call, triggering once per session or on every call.

Create a suggestion policy so agents are nudged toward `create_app` on every Vercel deploy call, without blocking the original tool.

> 📍 In Settings → Security Scanners → Overrides, open the relevant connector override page, add a suggestion policy targeting the Vercel deploy tool call, enter your preferred message referencing `create_app`, and set frequency to `Every call` or `Once per session`.

— Exact navigation, example, and trigger frequency given. 3.2.0 (2026-11-06)

---

04 **Tunnel connector observability: probes, metrics, and structured logs**

NEW 80

Tunnel connectors now support optional liveness and readiness probes to distinguish a responsive connector from one with an established tunnel session, an optional Prometheus-compatible metrics endpoint covering connector readiness, endpoint connectivity, and connection failures, and optional JSON logs with timestamp, severity, and message fields; text remains the default and now uses `key=value` fields.

— Names every surface: probes, metrics endpoint, log fields. 3.1.0 (2026-11-05)

---

05 **Direct artifact editing without invoking an agent** NEW

70

Edit doc, CSV/TSV, and HTML artifacts directly from the artifact page without invoking an agent; an agent's `edit_artifact` call that would overwrite a newer manual save is refused.

— Names formats and refusal mechanism, no UI steps. 3.1.0 (2026-11-05)

---

06 **Agent spend and run limit controls via MCP** NEW 60

Read an agent's spend and run limits, including current usage and reset times, and set or remove daily and weekly spend caps through the Runlayer Assistant and MCP clients.

Inspect or cap an agent's daily spend limit from an MCP client without leaving your automation workflow.

> 📍 In the Runlayer Assistant or any MCP client, ask: 'Show current spend and run limits for agent <agent-id>, then set a daily spend cap of $10.'

— Names caps and usage fields but no exact command. 3.2.0 (2026-11-06)

---

07 **In-place artifact title renaming** NEW 60

Editors can rename artifacts by double-clicking the title or editing it while `Edit text` is open; renaming does not create a new version, and a blank title falls back to the file

name.

— Exact UI interaction and versioning behavior given. 3.2.0 (2026-11-06)

08 **Reorganized Security Scanners settings** IMPROVED 60

Settings → Security Scanners now opens with tabs for Scanners, Security Response, Scan budgets, and Overrides, a scanner list grouped by category with a detail pane, and one page per connector or client override.

— Names all four tabs but no deeper mechanism. 3.2.0 (2026-11-06)

09 **Per-schedule and per-trigger model selection** NEW 60

Run scheduled tasks and connector triggers on a model different from the agent's default — pick a cheaper model for frequent checks and a stronger one for heavier reports — via the task or trigger dialog.

Assign a lightweight model to an hourly health-check schedule and a more capable model to a weekly report trigger, keeping costs down without changing the agent's default.

> 📍 In the task or trigger dialog, open the model selector, choose the model the agent's AI provider serves, and save — the schedule or trigger will use that model instead of the agent's default.

— Exact dialog steps given via example. 3.1.0 (2026-11-05)

THINNER COVERAGE BELOW

10 **Automatic benchmarking of procedure drafts** IMPROVED 55

Procedure drafts are now benchmarked automatically by replaying the agent runs they were built from, so a draft reaches approval-ready or goes live without a manual benchmark run; the review panel shows queued, running, or failed status.

— Explains mechanism and review states, no UI path. 3.1.0 (2026-11-05)

11 **Length caps on PII and policy regex patterns** IMPROVED 55

Custom PII rule patterns are now capped at 1,000 characters and policy regex conditions at 500 characters; existing longer patterns continue to work.

— Exact numeric caps given, no mechanism beyond caps. 3.1.0 (2026-11-05)

12 **Kubernetes Gateway API publishing via Runlayer operator** NEW 50

Runlayer instances can now be published using the Kubernetes Gateway API alongside or instead of Ingress, configured via the Runlayer operator.

– No mechanism beyond publish method; no config example given. 3.2.0 (2026-11-06)

13 **Explained and auditable reveal of redacted agent runs**

IMPROVED 50

Opening a teammate's redacted agent run in chat now explains why messages are hidden, and people with unredacted access can reveal messages inline, with each reveal recorded in the audit log.

– Explains mechanism and audit logging, no UI path. 3.2.0 (2026-11-06)

14 **Restricted access to Runlayer-managed agents** BREAKING

50

Access to Runlayer-managed agents such as the Runlayer Assistant is now restricted to admins who manage security policies; an agent's owner can no longer grant edit or View all runs access.

– Explains before/after access change, no UI path. 3.1.0 (2026-11-05)

15 **Admin team inspection without joining** NEW 40

Admins can open any team from the Teams page and inspect its members and shared assets without joining it.

– Thin: names UI page only, no mechanism. 3.2.0 (2026-11-06)

16 **Link from Organization to Personal API keys** IMPROVED

40

The Organization API keys page now links to Personal API keys so users can find account-tied keys without navigating through Settings.

– Thin: just a navigation link added. 3.1.0 (2026-11-05)

17 **Agent identity shown in session panel** IMPROVED 35

Sessions now show the agent's avatar and name instead of a generic Runlayer label, with a direct link to the agent from the session panel.

– Cosmetic change, no mechanism or config detail. 3.2.0 (2026-11-06)

└──▷ **BREAKING ON UPGRADE**

! The Kubernetes operator now requires at least 4096 MiB of memory per slow-pool worker on `components.topicCpu` (up from 2048 MiB); any `RunlayerInstance` with `topicCpu.gunicornWorkers` set to 3 or more without a matching `resources` memory override fails validation after upgrade and is not reconciled — including version bumps — until memory is raised or worker count is lowered.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Straiker

Feedback ⓘ

TrueFoundry & Straiker | AI Gateway Security for Agents (2026-11-05) BLOG ↗

RANK

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

> Straiker shipped a TrueFoundry AI Gateway integration that lets Straiker act as a Custom Guardrail inspecting LLM, agent, and MCP traffic at the gateway layer, backed by configurable enforcement strategies, expanded Discover AI asset mapping, Ascend AI adversarial testing, and flexible SaaS/on-premises deployment.

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

01 **Enforcement strategy and timeout controls** NEW 90

Guardrail enforcement behavior can be set to one of three strategies — `enforce` (block on error), `enforce_but_ignore_on_error` (fail-open), and `audit` (record only) — and response-wait time is tuned with the `STRAIKER_TIMEOUT_S` setting.

Set a 200 ms timeout and fail-open behavior for the Straiker guardrail so availability is preserved when the detection service is slow or unreachable.

env

```
STRAIKER_TIMEOUT_S=0.2
```

— Named flags plus a runnable config example TrueFoundry & Straiker | AI Gateway Sec…

02 **TrueFoundry AI Gateway integration setup** NEW 75

Straiker can be configured as a Custom Guardrail in TrueFoundry's AI Gateway → Guardrails to route model requests, agent interactions, and MCP tool calls through Straiker's detect API without an adapter or wrapper service. The connection is established from Straiker Defend by selecting TrueFoundry from the available gateway integrations, naming the connection, and supplying credentials.

Start observing live agent and MCP traffic through Straiker without blocking anything, before enabling enforcement policies.

📍 In Straiker Defend, go to Integrations, select TrueFoundry, name the connection, and provide credentials. In TrueFoundry, go to AI Gateway →

Guardrails, register Straiker as a Custom Guardrail, set enforcement strategy to `audit`, and attach it to your target model or guardrail policy.

— Names exact UI path and connection steps, no API payload shown

TrueFoundry & Straiker | AI Gateway Sec…

03 **Detect API verdicts across gateway enforcement points** NEW

73

Straiker's detect API returns three enforcement verdicts — Allow, Detect, and Block — applied by TrueFoundry within the gateway path so applications get runtime protection without code changes. Enforcement covers four points in the TrueFoundry gateway: LLM input, LLM output (non-streaming), MCP tool pre-invocation, and MCP tool post-invocation.

— Mechanism and scope named, but no command to invoke it

TrueFoundry & Straiker | AI Gateway Sec…

04 **SaaS and on-premises deployment with VPC isolation** NEW

70

Straiker can be deployed as SaaS or on-premises; with the self-hosted gateway, in-band traffic stays inside the customer's VPC and only detection metadata reaches the Straiker control plane, with data encrypted using TLS 1.2 or higher in transit and AES-256 at rest.

— Concrete security specs but no deployment steps provided

TrueFoundry & Straiker | AI Gateway Sec…

05 **Automated adversarial campaigns via Ascend AI** IMPROVED

60

Straiker's Ascend AI now runs automated multi-turn adversarial campaigns covering prompt injection, goal hijacking, tool manipulation, sensitive-data exfiltration, remote code execution, and lateral movement through connected systems.

— Lists attack categories but no setup or trigger mechanism given

TrueFoundry & Straiker | AI Gateway Sec…

THINNER COVERAGE BELOW

06 **Discover AI mapping of TrueFoundry agents** IMPROVED

50

Straiker's Discover AI capability now maps TrueFoundry-deployed agents, models, tools, MCP servers, permissions, and data-access paths into a security view of the agentic ecosystem.

— Describes scope of mapping but no concrete navigation or output format

TrueFoundry & Straiker | AI Gateway Sec…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 6, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue