THE CYBER TOOLCHAIN

NO. 082

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 22, 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 284 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 52 tools

▾ // THE BRIEF

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

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. Falco 0.45.0 can match raw bytes in rule conditions, so detections stop depending on printable-string

vim ? help

assumptions, and adds GPU resource hijacking rules. NodeZero's pentests now follow footholds beyond the original scope through tunnels and Azure Run Command. StackHawk re-authenticates mid-scan, so a dropped session no longer silently hollows out an active DAST run.

DETECT

Write rules against bytes that never render as text, and catch GPU theft

Falco 0.45.0 matches and emits raw bytes in conditions and output, which is what you need for obfuscated payloads, binary protocol framing and non-UTF8 argv that string comparison quietly misses; new rules cover GPU resource hijacking, a real cost in shared AI infrastructure. Tetragon's `matchUserCallers` selector filters uprobe events by the user-space callstack, so a hook on a common library function fires only when the interesting caller is the one reaching it.

Falco · tetragon

PROTECT

Stop a web scan from reporting clean because the session died halfway through

HawkScan 6.5.0 re-authenticates mid-scan when a session drops, closing the classic DAST failure where the second half of an authenticated crawl runs as an anonymous user and finds nothing; a rescan fix also stops legitimate headers being stripped from replayed requests. CrowdSec built AppSec into a full WAF with ModSecurity rule import and OWASP CRS collections, plus alpha bot detection using proof-of-work challenges and device fingerprinting.

StackHawk · CrowdSec

IDENTIFY

See which SaaS apps and AI agents your people actually signed up for

JumpCloud's SaaS Management discovers shadow IT and adds a Windows Device Patch Report with CSV export, so the unsanctioned-app conversation starts from a list rather than an expense report. Okta's agent platform pairs shadow AI discovery with runtime enforcement and ephemeral credentialing — the point being that agents currently run on long-lived human credentials nobody inventories.

JumpCloud · Okta for AI Agents

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Exploitation & C2

Dradis

Feedback ⓘ New in Dradis Pro v5.4 (2026-09-22) BLOG ↗ RANK

Pentest reporting for consultancies that never put a client

> Dradis Pro v5.4 adds QA states to gate client deliverables, an in-editor Roslin AI drafting tool for Content Blocks, new Webhook lifecycle events, dropdown-enforced List fields, and expanded custom field support in the Jira and DevOps integrations.

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

01 **QA states on Evidence gate report exports** NEW 70

Adds QA states to Evidence so teams can track review status and gate client-facing deliverables — Excel and Word exports now only include Evidence marked as published. In practice, a reviewer opens a piece of Evidence, sets its QA state to Published from the Evidence view, and only published items appear in the subsequent Word or Excel export.

Use QA states to gate Evidence out of client deliverables — only published Evidence appears in Excel and Word exports.

📍 In Dradis Pro, open a piece of Evidence, then set its QA state from the Evidence view (e.g. mark it Published once reviewed). Repeat for each Evidence item before exporting your Word or Excel report.

Mark reviewed Evidence as published so only approved findings appear in client-facing Excel and Word exports.

📍 1. Open a piece of Evidence in your project. 2. Locate the QA state control in the Evidence view. 3. Set the state to Published once review is complete. 4. Export the report — only Evidence marked as published will appear in the Excel or Word output.

Mark Evidence as reviewed before export to ensure only published Evidence reaches the client Word or Excel report.

📍 In a project, open a piece of Evidence › locate the QA state selector in the Evidence view › set the state to Published before generating a Word or Excel export.

Gate client deliverables by marking each piece of Evidence with a QA state before export, ensuring only reviewed Evidence appears in Word and Excel reports.

📍 In a project, open a piece of Evidence, then set its QA state to `Published` from the Evidence view. When you export the report (Excel or Word), only Evidence marked as published will be included.

Gate client deliverables by marking Evidence through QA review before export — only published Evidence appears in Word and Excel reports.

📍 In a project, open a piece of Evidence, then use the QA state selector in the Evidence view to mark it as reviewed/published before generating your Word or Excel export.

Mark Evidence as reviewed before export so only published Evidence appears in client-facing Word and Excel deliverables.

📍 In Dradis, open a piece of Evidence › locate the QA state control in the Evidence view › set the state to `Published` before generating the report.

After completing review of a piece of evidence, mark it as published so it is included in the Word or Excel report export — unpublished evidence is now excluded automatically.

📍 In the project, open a piece of Evidence › set its QA state to `Published` from the Evidence view › generate the Word or Excel export — only published Evidence will appear in the deliverable.

Track Evidence review status before sending a report to a client, keeping unreviewed material out of exported deliverables.

📍 In a project, open a piece of Evidence, then set its QA state from the Evidence view. Only Evidence marked as published will appear in Excel and Word exports.

— Explains mechanism and export scope but no config surface

`New in Dradis Pro v5.4 (2026-09-22)` `launch-20260922-b5f44454`

02 **Roslin AI drafting in Content Blocks** NEW 60

Adds a Roslin agent tab to Content Blocks to draft report sections (executive summaries, methodology, scope notes) from the full project content including Issues and Evidence, without leaving the block. A user opens a Content Block, selects the Roslin tab, and prompts it to draft a section based on the project's findings.

Draft an executive summary or methodology section using Roslin directly inside a Content Block, drawing on the full project context.

📍 Open a Content Block in your project, navigate to the Roslin tab, and ask it to draft the section — for example: 'Draft an executive summary based on all Issues

and Evidence in this project.'

Draft a recurring report section (e.g. executive summary) from full project content without leaving the Content Block.

📍 1. Open a Content Block in your project. 2. Select the Roslin tab. 3. Ask Roslin to draft the section — for example: 'Draft an executive summary based on this project'. 4. Roslin responds using the full project content, including Issues and Evidence.

Draft an executive summary for a Content Block using the full project context — Issues and Evidence included — without leaving the block.

📍 Open a Content Block › select the `Roslin` tab › type a prompt such as 'Draft an executive summary based on the findings in this project' › review and insert the generated text.

Draft an executive summary for a Content Block using Roslin, drawing on all Issues and Evidence already in the project rather than rewriting from scratch.

📍 Open a Content Block, navigate to the Roslin tab, and ask Roslin to draft the section (e.g. 'Draft an executive summary based on all Issues and Evidence in this project').

Draft a recurring report section such as an executive summary directly from all project findings without leaving the Content Block.

📍 Open a Content Block, navigate to the Roslin tab, and prompt Roslin to draft the section — it reads from the full set of project Issues and Evidence to produce the content.

Draft an executive summary or methodology section for a recurring Content Block using the full project context without leaving the block.

📍 In Dradis, open a Content Block › select the `Roslin` tab › type a prompt such as 'Draft an executive summary based on all Issues and Evidence in this project' › review and save the generated content.

Draft an executive summary for a recurring Content Block without leaving Dradis, using the full project context (Issues and Evidence) as source material.

📍 Open a Content Block › click the `Roslin` tab › type a prompt such as 'Draft an executive summary based on the issues and evidence in this project' › review and accept the drafted section.

Draft a recurring report section such as an executive summary or methodology note without leaving the Content Block.

> 📍 Open a Content Block, navigate to the Roslin tab, and ask Roslin to draft the section — it draws on the full project content including Issues and Evidence.

— Names UI path and inputs but no API detail `New in Dradis Pro v5.4 (2026-09-22)` `launch-20260922-b5f44454`

THINNER COVERAGE BELOW

03 **Webhook events for Evidence, Node, and Note** NEW 55

Adds Evidence, Node, and Note CRUD events to Webhooks, enabling external systems to react to evidence lifecycle changes.

— Names event types but no payload or endpoint detail `New in Dradis Pro v5.4 (2026-09-22)`

04 **Expanded custom field support in Jira and DevOps integrations**

IMPROVED 55

Adds support for float custom fields and multiselect custom fields when sending issues to Jira, and adds support for numeric (integer and double) Work Item fields and required fields in the DevOps integration.

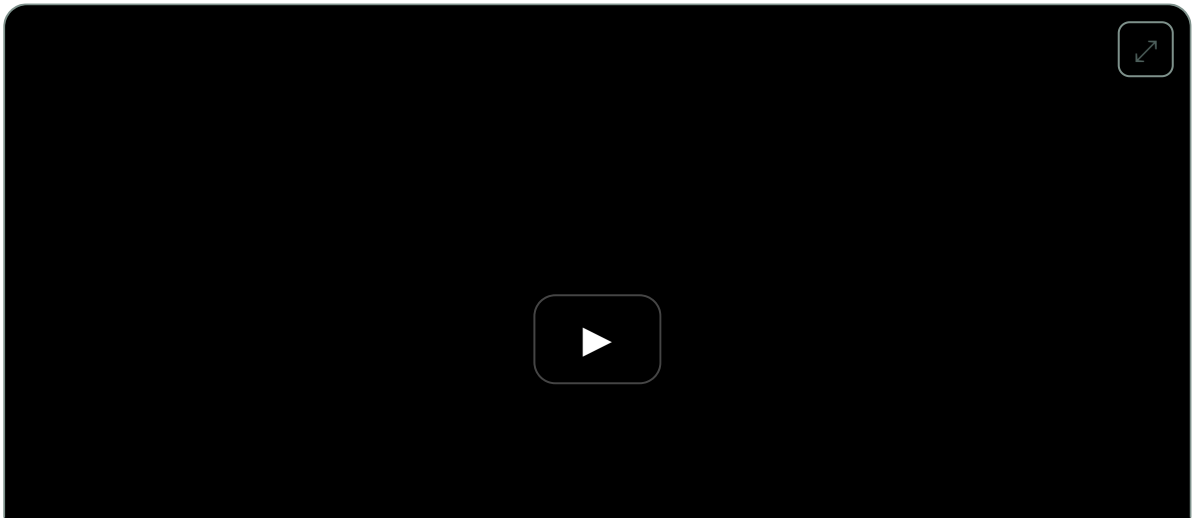
— Names field types and both integrations but no setup steps `New in Dradis Pro v5.4 (2026-09-22)`

05 **List fields rendered as dropdowns** IMPROVED 50

Renders List fields as selects (dropdowns) when editing existing Issues, Evidence, and Content Blocks, enforcing consistent field values across a project.

— States the change and scope but no usage steps `New in Dradis Pro v5.4 (2026-09-22)`

└──▷ ALSO FROM THESE RELEASES





New in Dradis Pro v5.4 (2026-09-22) BLOG ↗ LATEST STABLE

Dradis Pro v5.4 adds QA states to Evidence, Roslin AI drafting for Content Blocks, List fields as dropdowns, and Webhook CRUD events.

└──▷ **HOW TO FIND IT**

After completing review of a piece of evidence, mark it as published so it is included in the Word or Excel report export — unpublished evidence is now excluded automatically.

> 📍 In the project, open a piece of Evidence › set its QA state to `Published` from the Evidence view › generate the Word or Excel export — only published Evidence will appear in the deliverable.

Draft an executive summary for a recurring Content Block without leaving Dradis, using the full project context (Issues and Evidence) as source material.

> 📍 Open a Content Block › click the `Roslin` tab › type a prompt such as 'Draft an executive summary based on the issues and evidence in this project' › review and accept the drafted section.

Adds `Evidence`, `Node`, and `Note` CRUD events to Webhooks, enabling external systems to react to evidence lifecycle changes.

Adds QA states to Evidence so teams can track review status and gate client-facing deliverables — Excel and Word exports now only include published Evidence.

Adds a Roslin agent tab to Content Blocks to draft report sections (executive summaries, methodology, scope notes) from the full project content including Issues and Evidence.

Renders List fields as selects (dropdowns) when editing existing Issues, Evidence, and Content Blocks, enforcing consistent field values across a project.

Adds support for numeric (integer and double) Work Item fields and required fields in the DevOps integration.

+1 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Veracode

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

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

> Veracode expanded platform integrations and scanning accuracy this window, adding Jira 11 support to its Jira Server integration, refining Python scanner flaw matching against a baseline scan, and redesigning its EASM interface.

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

01 **Jira Server integration updated for Jira 11** IMPROVED 60

Veracode Integration for Jira Server `v5.0.2` adds support for Jira 11, including Java 21 and platform compatibility with Jira 11.3.10.

– Names exact version and compatibility targets but no setup steps.

changelog-20260922-9065506b

THINNER COVERAGE BELOW

02 **Python scanner flaw matching baseline** BREAKING 50

The enhanced Python scanner now uses each application's initial scan baseline for consistent flaw matching across subsequent scans. Customers who transitioned to the enhanced Python scanner on July 22, 2026 may find previously identified flaws no longer match historical findings, and previously applied mitigations may need to be reapplied.

– Explains mechanism and migration impact but no concrete remediation steps.

changelog-20260922-9065506b

03 **Redesigned EASM UI** IMPROVED 10

Veracode shipped a redesigned UI for its external attack surface management (EASM) capability.

– Mentioned only in passing with no detail on what changed. changelog-20260922-9065506b

└──▷ **BREAKING ON UPGRADE**

! Customers who transitioned to the enhanced Python scanner on July 22, 2026 may find that previously identified flaws no longer match historical findings and previously applied mitigations may need to be reapplied.

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

Veracode adds Jira 11 support, a redesigned EASM UI, and improved Python scanner flaw matching.

Veracode Integration for Jira Server v5.0.2 adds support for Jira 11, including Java 21 and platform compatibility with Jira 11.3.10.

Enhanced Python scanner now uses each application's initial scan baseline for consistent flaw matching across subsequent scans.

└──▷ **BREAKING ON UPGRADE**

! Customers who transitioned to the enhanced Python scanner on July 22, 2026 may find that previously identified flaws no longer match historical findings and previously applied mitigations may need to be reapplied.

WAS THIS USEFUL? 👍 👎

StackHawk

Feedback ⓘ snapshot-20260922 NOTES ↗ RANK

StackHawk is an application security tool built for developers. Automated scanning for the entire dev pipeline combined with continuous documentation and remediation.

> HawkScan 6.5.0 adds mid-scan session re-authentication to catch dropped sessions during active scans, a one-line versioned CLI installer, and org-wide secret resolution, alongside a rescan fix that stops legitimate headers from being stripped from replayed requests.

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

01 **Mid-scan session verification** NEW 91

Adds an opt-in `sessionVerification` block under `app.authentication` in `stackhawk.yml` with `pollFrequency` (default 15), `activeScanPollFrequency` (default 5), `maxUnauthenticatedRequestPercent` (default 5), and `disabled` fields to actively poll `authentication.testPath` and re-authenticate sessions mid-scan rather than relying on passive response indicators. Supports both HTTP `testPath` and gRPC `testPath.grpcTestMethod` styles, with `testPath.type` values `HEADER` and `BODY` honored, and surfaces a coverage warning at scan completion if too many mid-scan checks fail unauthenticated, catching scans that silently ran as an anonymous user; `disabled: true` turns off active checks while still requiring an indicator-based check.

Enable mid-scan session re-authentication so HawkScan detects and recovers from dropped sessions during active attack phases, with a warning if more than 10% of checks are unauthenticated.

yaml

```yaml
app:
  authentication:
    sessionVerification:
      pollFrequency: 50
      activeScanPollFrequency: 10
      maxUnauthenticatedRequestPercent: 10
      disabled: false
```

Enable mid-scan session re-authentication so unauthenticated drift during an active attack scan is caught and corrected automatically, with tighter polling during active scanning.

```yaml
app:
  authentication:
    sessionVerification:
      pollFrequency: 20
      activeScanPollFrequency: 3
      maxUnauthenticatedRequestPercent: 5
```

Enable mid-scan session verification for an API scan that frequently loses its authenticated state during active scanning.


```yaml
app:
  authentication:
    sessionVerification:
      pollFrequency: 20
      activeScanPollFrequency: 3
      maxUnauthenticatedRequestPercent: 5
```


— Full config block, fields, defaults and behavior all named `snapshot-20260922`

02 **One-line HawkScan CLI installer** NEW 88

Adds a one-line installer script, `curl -sSf https://download.stackhawk.com/hawk/install.sh | sh`, that detects OS and architecture and verifies SHA-256 against the published manifest. Supports `HAWK_INSTALL_VERSION` to pin the CLI version (recommended for reproducible CI pipelines) and `HAWK_INSTALL_DIR` to control install location, which defaults to `~/.hawk/bin`.

Pin the hawk CLI to a specific version in CI so your pipeline stays reproducible and does not break when a new release drops.

```
$ HAWK_INSTALL_VERSION=6.5.0 HAWK_INSTALL_DIR=/usr/local/bin curl -sSf https://download.stackhawk.com/hawk/install.sh | sh
```

Pin the HawkScan CLI to a specific version in CI for a reproducible pipeline, installing to a custom directory.

```
$ HAWK_INSTALL_VERSION=6.5.0 HAWK_INSTALL_DIR=/opt/hawk/bin curl -sSf https://download.stackhawk.com/hawk/install.sh | sh
```

— Runnable install command with named env vars and defaults snapshot-20260922

THINNER COVERAGE BELOW

03 **Headers no longer stripped on rescan** IMPROVED 55

Entity and transport headers — Content-Type, Content-Length, Accept, Accept-Encoding, Host, User-Agent, Referer — are no longer stripped from replayed requests during rescans, preventing spurious 415 responses from masking unresolved findings.

— Names headers and effect but no action step for reader product docs

04 **Organization-level secret resolution** NEW 50

Adds organization-level secret resolution so `${VARIABLE}` references in scan configuration now resolve against org-shared secrets, not only per-user secrets; personal secrets still take precedence, and unresolved names now emit a named warning instead of silently resolving to an empty value.

— Explains mechanism but no concrete config path shown snapshot-20260922

05 **CLI minimum-version enforcement messaging** IMPROVED 45

Adds clear error messaging when the HawkScan CLI version is below the platform-enforced minimum, reporting the minimum supported version and upgrade location and

exiting non-zero instead of emitting a generic HTTP error; deprecated-but-not-blocked versions now log an advisory once per scan rather than on every platform call.

— Describes behavior change but no exact command or code `snapshot-20260922`

`snapshot-20260922` **NOTES ↗** **LATEST** **STABLE**

HawkScan 6.5.0 adds mid-scan session verification, org-wide secrets, and a one-line CLI installer with version pinning.

└──▷ **TRY IT**

Pin the hawk CLI to a specific version in CI so your pipeline stays reproducible and does not break when a new release drops.

```
$ HAWK_INSTALL_VERSION=6.5.0 HAWK_INSTALL_DIR=/usr/local/bin curl -sSf https://download.stackhawk.com/hawk/install.sh | sh
```

Enable mid-scan session re-authentication so unauthenticated drift during an active attack scan is caught and corrected automatically, with tighter polling during active scanning.

stackhawk.yml

```
app:
  authentication:
    sessionVerification:
      pollFrequency: 20
      activeScanPollFrequency: 3
      maxUnauthenticatedRequestPercent: 5
```

Adds `app.authentication.sessionVerification` config block to re-check and re-authenticate sessions mid-scan by polling `authentication.testPath`, with tunable `pollFrequency` (default 15), `activeScanPollFrequency` (default 5), and `maxUnauthenticatedRequestPercent` (default 5); a scan that spent too much time logged out now finishes with a coverage warning instead of a silently thin result.

Adds `disabled: true` field inside `app.authentication.sessionVerification` to explicitly turn off session verification while still requiring an indicator-based check.

Adds a one-line installer script at `curl -sSf https://download.stackhawk.com/hawk/install.sh | sh` that detects OS and architecture, verifies SHA-256 against the published manifest, and supports `HAWK_INSTALL_VERSION` and `HAWK_INSTALL_DIR` environment variables (default install location `~/.hawk/bin`).

Adds `HAWK_INSTALL_VERSION` environment variable to pin the `hawk` CLI version during install — recommended for reproducible CI pipelines.

> Adds `HAWK_INSTALL_DIR` environment variable to control where the `hawk` binary is installed (default `~/.hawk/bin`).
>
> +2 more

WAS THIS USEFUL?  

GitHub Feedback ⓘ

Refreshed repository pull requests page generally available (2026-09-22) BLOG ↗

RANK

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

> GitHub shipped a generally-available refresh of the repository pull requests page with advanced search and bulk actions, added a full credential inventory export for Enterprise, and rolled out the Grok 4.7 model in Copilot.

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

01 Enterprise credential inventory export NEW

87

GitHub Enterprise adds a new View enterprise credentials fine-grained permission granting members access to the credential inventory alongside enterprise owners, plus new paginated REST API endpoints to programmatically retrieve the complete inventory. An Export CSV option in enterprise Settings -> Authentication Security -> Credentials downloads the full inventory, which includes owners, scopes and permissions, creation and expiration dates, last-used dates, and target organizations or repositories; it can be filtered by user, app, credential type, or organization, and covers SSH keys, classic and fine-grained personal access tokens, OAuth App access tokens, and GitHub App user-to-server and installation tokens.

Export a one-off CSV snapshot of all enterprise credentials for offline review or audit when you don't need automation.

> 📍 In your enterprise, go to Settings -> Authentication Security -> Credentials and click Export CSV next to the Overview section.

— Names permission, UI path and data covered but not exact endpoint paths.

GitHub Enterprise adds credential inven…

02 Refreshed pull requests page generally available NEW

85

The pull requests page now supports advanced search with `AND` and `OR` keywords and nested searches, and bulk actions to close, label, or change milestones on one or more pull requests at once. A collapsible sidebar gives quick access to common filters like Authored by me and Involves me and can be collapsed to expand the list view; a compact presentation mode fits more pull requests on the page. The list now surfaces status check counts, stack indicators, unread-update indicators, milestones, linked

issues, and author badges directly, clicking a review status filters the list by that status, and full profile names are shown when enabled for an enterprise.

Narrow a busy repository's PR list to open, approved PRs that still need a specific team's attention — combining OR and AND in one query.

> 📍 In the repository pull requests page, enter into the search bar: is:open review:approved AND (involves:@my-org/security-team OR involves:@my-org/platform-team)

Quickly close or re-label a batch of stale PRs without opening each one individually.

> 📍 In the repository pull requests page, check the checkboxes next to the target pull requests, then open the `Label` or `Close` bulk-action menu that appears at the top of the list.

— Names exact search syntax and UI mechanics, with runnable query example.

`Refreshed repository pull requests page…`

---

03 **Grok 4.7 model in Copilot** NEW 65

The Grok 4.7 reasoning model is now available in GitHub Copilot across VS Code, JetBrains, Xcode, Eclipse, CLI, and the cloud agent. Copilot Enterprise and Business administrators manage access via the model policy in Copilot settings, where new models are auto-enabled under the global default unless explicitly disabled.

Restrict or permit Grok 4.7 for your organisation by toggling its model policy in Copilot settings.

> 📍 In GitHub, go to your organisation's Settings › Copilot › Policies, find `Grok 4.7` under model enablement, and set it to `Enabled` or `Disabled` as required.

— Names surfaces and admin toggle path but little mechanism detail.

`Grok 4.7 is now available in GitHub Cop…`

`Refreshed repository pull requests page generally available (2026-09-22)`

BLOG ↗ LATEST STABLE

GitHub's refreshed pull requests page is now GA, with AND/OR advanced search, bulk actions, and a collapsible sidebar.

└──▷ **HOW TO FIND IT**

Narrow a busy repository's PR list to open, approved PRs that still need a specific team's attention — combining OR and AND in one query.

> 📍 In the repository pull requests page, enter into the search bar: is:open review:approved AND (involves:@my-org/security-team OR involves:@my-

org/platform-team)

Quickly close or re-label a batch of stale PRs without opening each one individually.

> 📍 In the repository pull requests page, check the checkboxes next to the target pull requests, then open the `Label` or `Close` bulk-action menu that appears at the top of the list.

Advanced search on the pull requests page now supports `AND` and `OR` keywords as well as nested searches.

Bulk actions let you close, label, or change milestones on one or more pull requests at once.

Collapsible sidebar provides quick access to common filters such as `Authored by me` and `Involves me`, and can be collapsed to expand the list view.

Compact presentation mode fits more pull requests on the page.

Pull request list now surfaces status check counts, stack indicators, and unread-update indicators.

`+3 more`

`GitHub Enterprise adds credential inventory exports (2026-09-21)` **BLOG ↗**

STABLE

GitHub Enterprise gains a full credential inventory export — CSV from the UI or via new paginated REST API endpoints.

└──▷ HOW TO FIND IT

Export a one-off CSV snapshot of all enterprise credentials for offline review or audit when you don't need automation.

> 📍 In your enterprise, go to Settings -> Authentication Security -> Credentials and click `Export CSV` next to the `Overview` section.

Adds a new `View enterprise credentials` fine-grained permission that grants members access to the enterprise credential inventory alongside enterprise owners.

Adds new paginated REST API endpoints to programmatically retrieve the complete enterprise credential inventory for reporting and automation.

Adds an `Export CSV` option in enterprise Settings -> Authentication Security -> Credentials to download the full credential inventory.

Surfaces credential metadata in the inventory including owners, scopes and permissions, creation and expiration dates, last-used dates, and target organizations or repositories.

Supports filtering the credential inventory by user, app, credential type, or organization.

+1 more

Grok 4.7 is now available in GitHub Copilot (2026-09-21) BLOG ↗ STABLE

Grok 4.7 reasoning model now available in GitHub Copilot across VS Code, JetBrains, Xcode, Eclipse, CLI, and cloud agent.

└──▷ **HOW TO FIND IT**

Restrict or permit Grok 4.7 for your organisation by toggling its model policy in Copilot settings.

> 📍 In GitHub, go to your organisation's Settings › Copilot › Policies, find `Grok 4.7` under model enablement, and set it to `Enabled` or `Disabled` as required.

Copilot Enterprise and Business administrators can manage access to Grok 4.7 via the model policy in Copilot settings, with new models auto-enabled under the global default unless explicitly disabled.

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ announced · Sep 21, 2026

BLOG ↗ RANK

SonarQube is the independent AI code verification platform trusted by 7M+ developers & organizations to help build secure, reliable software with confidence.

> SonarQube launched Hunter Agent, a new AI-driven verification layer that hunts for business-logic, access-control, and authentication flaws beyond the reach of static analysis and SCA, and wires the results straight into existing triage and compliance workflows.

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

01 **Hunter Agent AI vulnerability detection** NEW 73

Hunter Agent is a new AI-driven verification layer that runs structured multi-stage playbooks to detect business-logic, broken access control, and authentication/session vulnerabilities — classes of flaw that static analysis and SCA miss. Findings surface as standard SonarQube issues (with severity, file, and line) inside the existing triage, assignment, and tracking workflow, and are also summarized on a dedicated Hunter Agent page; it can be run manually or on a daily schedule against long-lived branches. It produces a reproducible, auditable evidence trail from detection through remediation, mapped to DORA, NIS2, PCI DSS 4.0.1, HIPAA, GDPR, EU CRA, NYDFS, and FDA 21 CFR Part 11 requirements.

— Names mechanism, workflow integration, scheduling, and compliance mappings; no config/API surface given.

launch-20260921-74341ed0

> announced · Sep 21, 2026 BLOG ↗
>
> SonarQube Hunter Agent adds AI-driven detection of business logic, access control, and authentication flaws as standard issues.
>
> Adds Hunter Agent, a new AI-driven verification layer that detects business-logic, broken access control, and authentication/session vulnerabilities through structured multi-stage playbooks — classes of flaw that static analysis and SCA miss.
>
> Surfaces Hunter Agent findings as standard SonarQube issues (with severity, file, and line) in the existing triage, assignment, and tracking workflow, with results also summarized on a dedicated Hunter Agent page.
>
> Supports running Hunter Agent manually or on a daily schedule against long-lived branches to suit different team cadences.
>
> Produces a reproducible, auditable evidence trail from detection through triage and remediation, mapped to DORA, NIS2, PCI DSS 4.0.1, HIPAA, GDPR, EU CRA, NYDFS,

and FDA 21 CFR Part 11 requirements.

WAS THIS USEFUL? 👍 👎

Pi

Feedback ⓘ announced · Sep 19, 2026 BLOG ↗ RANK

Pi Security triages vulnerability and bug-bounty reports with agentic AI, analyzes the organisation's code and cloud context and generates fixes for developers.

> Pi Security launched an integration with Anthropic's Compliance API and introduced Sloane, an in-session agent that surfaces security context to developers as they write code.

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

01 **Sloane in-session security agent** NEW 45

Sloane is Pi's agent running natively in the coding session, surfacing security context to developers as code is written rather than in post-review.

— Explains what it does and when, but no usage detail. launch-20260919-f1816728

02 **Anthropic Compliance API integration** NEW 40

Pi Security integrates with Anthropic's Compliance API to maintain a complete compliance record of the security context delivered to developers.

— Names the integration but no mechanism or setup steps. launch-20260919-f1816728

> announced · Sep 19, 2026 BLOG ↗
>
> Pi Security integrates with Anthropic's Compliance API and delivers security context in-session via its Sloane agent.
>
> Integrates with Anthropic's Compliance API to maintain a complete compliance record of security context delivered to developers.
>
> Introduces Sloane, Pi's agent running natively in the coding session, to surface security context to developers as code is written rather than in post-review.

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ

Apps & PATs Inventory for GitHub Enterprise Server (2026-09-18) BLOG ↗ RANK

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

> StepSecurity extended its Apps & PATs inventory to GitHub Enterprise Server and added CC recipients, per-token-class schedules, and customizable email content to PAT expiry reminders.

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

01 PAT reminder CC recipients, custom schedules and email content

IMPROVED 93

PAT Governance reminders now support CC recipients so additional addresses (e.g. a security distribution list) are copied on every pre-expiry PAT reminder. Per-token-class reminder schedules add optional `classic` and `fine-grained` overrides that replace the shared default schedule (30, 7, 1 days before expiry) for one class. Per-token-class and per-reminder-band subject and body editors are prefilled with the default email and support placeholders for owner, token name, organization, days left, expiry date, and the management URL.

Route all PAT expiry reminders to a security team inbox in addition to the org notification address, and set a tighter fine-grained-only schedule.

📍 In the StepSecurity console, go to Settings › Control Evaluation › Reminders & Automation (PAT Governance Controls). Under CC Recipients, add your security distribution list address. Under Reminder Schedule, enable the fine-grained override and set custom days (e.g. 14, 3, 1) to replace the shared 30/7/1-day default for fine-grained tokens only.

Customize the 1-day pre-expiry email body for classic tokens to include your internal rotation runbook URL.

📍 In the StepSecurity console, go to Settings › Control Evaluation › Reminders & Automation (PAT Governance Controls). Select the classic token class and the 1 day reminder band, then edit the subject and body fields — use placeholders such as {{owner}}, {{token_name}}, {{organization}}, {{days_left}}, {{expiry_date}}, and {{management_url}} — and save. Fields left unchanged keep the built-in email.

— Names exact settings path, default schedule values, and placeholder fields.

PAT Governance: Reminder Recipients and…

THINNER COVERAGE BELOW

02 Apps & PATs inventory for GitHub Enterprise Server NEW

45

The Apps & PATs page now covers GitHub Enterprise Server organizations, inventorying installed GitHub Apps and fine-grained personal access tokens with permission, access scope, and status detail.

— Describes scope and data shown but no UI path or config given.

Apps & PATs Inventory for GitHub Enterp…

Apps & PATs Inventory for GitHub Enterprise Server (2026-09-18) BLOG ↗

STABLE

StepSecurity Apps & PATs Inventory now available for GitHub Enterprise Server organizations.

PAT Governance: Reminder Recipients and Custom Email Content (2026-09-18)

BLOG ↗ LATEST STABLE

StepSecurity PAT Governance now supports CC recipients and per-class, per-band custom email content for pre-expiry token reminders.

└──▷ HOW TO FIND IT

Route all PAT expiry reminders to a security team inbox in addition to the org notification address, and set a tighter fine-grained-only schedule.

📍 In the StepSecurity console, go to Settings › Control Evaluation › Reminders & Automation (PAT Governance Controls). Under CC Recipients, add your security distribution list address. Under Reminder Schedule, enable the fine-grained override and set custom days (e.g. 14, 3, 1) to replace the shared 30/7/1-day default for fine-grained tokens only.

Customize the 1-day pre-expiry email body for classic tokens to include your internal rotation runbook URL.

📍 In the StepSecurity console, go to Settings › Control Evaluation › Reminders & Automation (PAT Governance Controls). Select the classic token class and the 1 day reminder band, then edit the subject and body fields — use placeholders such as {{owner}}, {{token_name}}, {{organization}}, {{days_left}}, {{expiry_date}}, and {{management_url}} — and save. Fields left unchanged keep the built-in email.

Adds CC recipient support so additional addresses (e.g. a security distribution list) are copied on every pre-expiry PAT reminder sent by the organization.

> Adds per-token-class reminder schedules with optional `classic` and `fine-grained` overrides that replace the shared default schedule (30, 7, 1 days before expiry) for one class.
>
> Adds per-token-class and per-reminder-band subject and body editors, prefilled with the default email, supporting placeholders for owner, token name, organization, days left, expiry date, and the management URL.

WAS THIS USEFUL?  

JFrog

Feedback ⓘ announced · Sep 16, 2026 BLOG ↗ RANK

JFrog provides DevOps platform for software delivery, managing artifacts, dependencies, and continuous integration/deployment pipelines.

> JFrog Artifactory adds native LuaRocks repository support, letting teams host, proxy, and cache Lua packages and Kong plugins through local, remote, and virtual repositories.

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

01 **LuaRocks repository support in Artifactory** NEW 75

Artifactory adds LuaRocks as a native repository type, supporting local, remote, and virtual repository configurations for Lua packages. Remote repositories proxy and cache the public luarocks.org registry, serving cached rocks on subsequent requests. Local repositories let CI jobs publish `.rock` files and host private rocks and Kong plugins so Kong nodes can install versioned plugins from a single internal endpoint, while a virtual repository unifies private and cached public rocks behind a single URL for the `luarocks` client.

— Names repo types and client/file formats but no exact config paths.

launch-20260916-e6e0e2f5

> announced · Sep 16, 2026 BLOG ↗
>
> JFrog Artifactory adds native LuaRocks repository support for hosting, proxying, and caching Lua rocks.
>
> Adds LuaRocks as a native repository type in JFrog Artifactory, supporting local, remote, and virtual repository configurations for Lua packages.
>
> Supports proxying and caching the public luarocks.org registry through a remote repository, serving cached rocks on subsequent requests.
>
> Enables hosting private rocks and Kong plugins in a local LuaRocks repository, allowing CI jobs to publish `.rock` files and Kong nodes to install versioned plugins from a single internal endpoint.
>
> Provides a virtual LuaRocks repository that unifies private and cached public rocks behind a single URL for the `luarocks` client.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

Detection Engine Updates Version 2.172 (2026-09-14) BLOG ↗ RANK

GitGuardian Detection Engine scans code repositories to identify and alert on exposed secrets, credentials, and sensitive data.

> GitGuardian's Detection Engine 2.172 release substantially widens secret coverage, adding 10 new detectors, 2 new checkers, and 8 new secret analyzers, while also improving match recall across several existing detectors and retiring a deprecated Figma scope.

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

01 **Expanded detection recall for Scaleway, Splunk, Salesforce, and generic detectors**

IMPROVED 71

The Scaleway Token detector now catches `SCW_SECRET_KEY` and `SCALEWAY_SECRET_KEY` assignments, including documents where the word `scaleway` is absent. The Splunk Authentication Token detector now also matches an optional Splunk host alongside the token. The Salesforce OAuth2 Keys detector now captures the Salesforce instance host and validates against that instance's OAuth token endpoint (production or sandbox), returning `failed_to_check` when no host is found. The Username Password detector now also matches Spanish, Portuguese, French, and Turkish username/password assignment keywords. The `generic_high_entropy_secret` and `generic_encryption_key` detectors now report secrets/keys located near a localhost mention. The `generic_database_assignment_attached_port` detector has been merged into the `generic_database_assignment` detector. The Google API Key detector now catches keys containing only one digit.

— Names every affected detector, env vars and merged detectors with mechanism

Detection Engine Updates Version 2.172 …

02 **Ten new secret detectors added** NEW 61

Adds detectors for Brightcove OAuth credentials, Infor ION API Gateway OAuth client credentials (capturing the authorization host and token endpoint when present), generic Sanctum tokens, Telnyx API keys, Bifrost virtual keys (matched alone or together with the Bifrost deployment URL), Bitbucket HTTP access tokens with the `BBDC-` prefix, Jira Data Center personal access tokens, MetaApi cloud tokens, Appwrite instance encryption keys, and APIclub API keys.

— Names ten specific detectors and some matching mechanisms, but detection is passive

Detection Engine Updates Version 2.172 …

THINNER COVERAGE BELOW

03 **Eight new secret analyzers added** NEW 56

Adds secret analyzers for AWS Cognito OAuth 2.0 credentials with host, Alibaba Cloud IAM keys, Databricks access tokens with hostname, PyPI API tokens, Contentful Content Management API keys, DigitalOcean personal access tokens v1, DigitalOcean tokens, and DigitalOcean OAuth application tokens v1.

— Lists all eight analyzer targets but no operable interface for users

Detection Engine Updates Version 2.172 …

04 **Figma token analyzer scope updated** DEPRECATED 48

The Figma personal access token analyzer now reports the `folders:read` scope, replacing the deprecated `projects:read` scope.

— Names old and new scope but is a single small change

Detection Engine Updates Version 2.172 …

05 **New checkers for Artifactory and Splunk tokens** NEW 44

Adds a checker for Artifactory reference tokens with custom host support, and a checker for Splunk authentication tokens with custom host support.

— Names both checkers and host-support capability, no user-facing action

Detection Engine Updates Version 2.172 …

> Detection Engine Updates Version 2.172 (2026-09-14) BLOG ↗ LATEST STABLE
>
> GitGuardian Detection Engine 2.172 adds 10 new detectors, 8 new secret analyzers, and expanded recall for Scaleway, Salesforce, Splunk, and more.
>
> Adds a new detector for Brightcove OAuth credentials.
>
> Adds a new detector for Infor ION API Gateway OAuth client credentials, capturing the authorization host and token endpoint when present.
>
> Adds a new detector for generic Sanctum tokens.
>
> Adds a new detector for Telnyx API keys with the `KEY` prefix.
>
> Adds a new detector for Bifrost virtual keys, matched alone or together with the Bifrost deployment URL.
>
> +24 more

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Capstone

Feedback ⓘ 6.0.0-Alpha11 NOTES ↗ RANK

Capstone disassembly/disassembler framework for ARM, ARM64 (ARMv8), Alpha, BPF, Ethereum VM, HPPA, LoongArch, M68K, M680X, Mips, MOS65XX, PPC, RISC-V(rv32G/rv64G), SH, Sparc, SystemZ, TMS320C64X, TriCore, Webassembly, XCore and X86.

> Capstone 6.0.0-Alpha11 adds x86 vendor-specific disassembly modes, gates Xtensa ESP32-S3/HIFI3 instructions behind a new mode flag, and adds PowerPC mcrxr instruction support.

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

01 x86 vendor mode flags for AMD/Intel NEW 85

Adds `CS_MODE_X86_AMD` and `CS_MODE_X86_INTEL` mode flags for x86 to enable vendor-specific disassembly behavior, currently affecting `JO` / `JNO` operand size and notation in 64-bit mode; default x86 behavior is unchanged.

Disassemble 64-bit x86 code using Intel vendor semantics so that JO/JNO operands are decoded according to Intel conventions.

python

```
csh = Cs(CS_ARCH_X86, CS_MODE_64 | CS_MODE_X86_INTEL)
for insn in csh.disasm(code, 0x1000):
    print('0x%x:\t%s\t%s' % (insn.address, insn.mnemonic,
insn.op_str))
```

Disassemble 64-bit x86 code using AMD-style JO/JNO operand notation instead of the default representation.

python

```
cs = Cs(CS_ARCH_X86, CS_MODE_64 | CS_MODE_X86_AMD)
for i in cs.disasm(code, 0x1000):
    print(i.mnemonic, i.op_str)
```

— Names exact flags, affected instructions, and runnable usage. 6.0.0-Alpha11

launch-20260921-37c5428c

02 Xtensa ESP32-S3/HIFI3 mode gating NEW 80

Gates Xtensa ESP32-S3 and HIFI3 instructions behind the new `CS_MODE_XTENSA_ESP32S3` mode flag to avoid unintended decode on other Xtensa targets.

Disassemble Xtensa firmware built for the ESP32-S3 to correctly decode ESP32-S3- and HIFI3-specific instructions that are otherwise gated off.

```python
csh = Cs(CS_ARCH_XTENSA, CS_MODE_XTENSA_ESP32S3)
for insn in csh.disasm(firmware_bytes, 0x40000000):
    print('0x%x:\t%s\t%s' % (insn.address, insn.mnemonic,
insn.op_str))
```

Disassemble Xtensa firmware built for the ESP32-S3 to expose ESP32-S3 and HIFI3 instructions that are hidden in the default mode.

```python
cs = Cs(CS_ARCH_XTENSA, CS_MODE_XTENSA_ESP32S3)
for i in cs.disasm(firmware_bytes, 0x40000000):
    print(i.mnemonic, i.op_str)
```

— Names flag and its purpose with runnable example. `6.0.0-Alpha11` `launch-20260921-37c5428c`

THINNER COVERAGE BELOW

03 **PowerPC mcrxr instruction support** NEW 30

Adds support for the `mcrxr` instruction in the PowerPC disassembly module.

— Names the instruction but no example or mechanism. `launch-20260921-37c5428c`

`6.0.0-Alpha11` NOTES ↗ PRE-RELEASE

Capstone 6.0.0-Alpha11 adds `CS_MODE_X86_AMD` /INTEL x86 vendor modes, `CS_MODE_XTENSA_ESP32S3` gating, and PowerPC mcrxr instruction support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 6.0.0-Alpha11 https://github.com/capstone-
engine/capstone.git
# already have the repo? check out this version:
$ git checkout 6.0.0-Alpha11
```

└──▷ USE IT

Disassemble 64-bit x86 code using Intel vendor semantics so that JO/JNO operands are decoded according to Intel conventions.

```python
csh = Cs(CS_ARCH_X86, CS_MODE_64 | CS_MODE_X86_INTEL)
for insn in csh.disasm(code, 0x1000):
    print('0x%x:\t%s\t%s' % (insn.address, insn.mnemonic,
insn.op_str))
```

Disassemble Xtensa firmware built for the ESP32-S3 to correctly decode ESP32-S3- and HIFI3-specific instructions that are otherwise gated off.

```python
csh = Cs(CS_ARCH_XTENSA, CS_MODE_XTENSA_ESP32S3)
for insn in csh.disasm(firmware_bytes, 0x40000000):
    print('0x%x:\t%s\t%s' % (insn.address, insn.mnemonic,
insn.op_str))
```

Adds `CS_MODE_X86_AMD` and `CS_MODE_X86_INTEL` mode flags for x86 to enable vendor-specific disassembly behavior (currently affecting `JO` / `JNO` operand size in 64-bit mode).

Gates Xtensa ESP32-S3 and HIFI3 instructions behind the new `CS_MODE_XTENSA_ESP32S3` mode flag to avoid unintended decode on other Xtensa targets.

announced · Sep 21, 2026 BLOG ↗

Capstone: Adds `CS_MODE_X86_AMD` and `CS_MODE_X86_INTEL` mode flags to select alternative x86 disassembly for JO/JNO operands in 64-bit mode; default x86 behavior is unchanged.

└──▷ **USE IT**

Disassemble 64-bit x86 code using AMD-style JO/JNO operand notation instead of the default representation.

```python
cs = Cs(CS_ARCH_X86, CS_MODE_64 | CS_MODE_X86_AMD)
for i in cs.disasm(code, 0x1000):
    print(i.mnemonic, i.op_str)
```

Disassemble Xtensa firmware built for the ESP32-S3 to expose ESP32-S3 and HIFI3 instructions that are hidden in the default mode.

```python
cs = Cs(CS_ARCH_XTENSA, CS_MODE_XTENSA_ESP32S3)
for i in cs.disasm(firmware_bytes, 0x40000000):
    print(i.mnemonic, i.op_str)
```

Adds support for the `mcrxr` instruction in the PowerPC disassembly module.

WAS THIS USEFUL? 👍 👎

EMBA

Feedback ⓘ v2.0.4-summer_edition NOTES ↗ RANK

EMBA - The firmware security analyzer

> EMBA v2.0.4 migrates firmware extraction to the actively maintained binwalk-ng fork, makes the local connector endpoint configurable, and completes CSV logging across all modules.

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

01 Configurable local connector endpoint and TCP settings

IMPROVED 65

The local connector endpoint is now configurable, a change introduced via `PR #2069`, and `MIN_TCP_SERV` has been moved to defaults where it can be adjusted via scanning profiles.

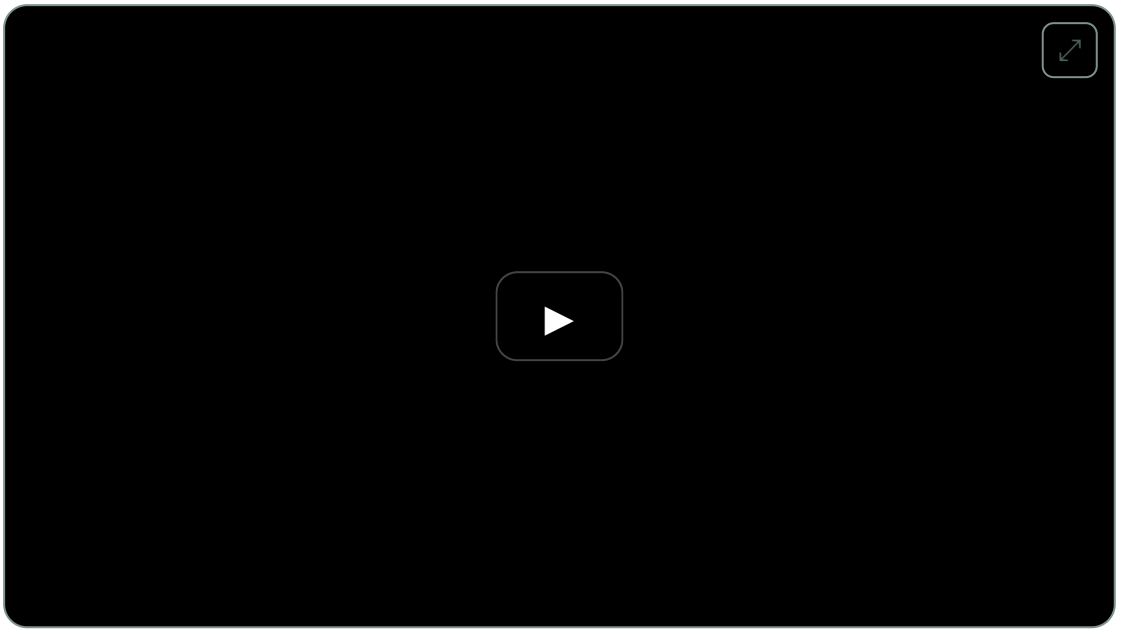
— Names PR and config key but no further mechanism v2.0.4-summer_edition

THINNER COVERAGE BELOW

02 Migration to binwalk-ng for firmware extraction IMPROVED

45

EMBA migrates firmware extraction from binwalk to the actively maintained `binwalk-ng` fork.



— Names the fork but no usage or migration detail v2.0.4-summer_edition

03 CSV logging across all modules IMPROVED

35

Completes CSV output logging across all modules for structured, machine-readable results.

— States scope but no format specifics beyond CSV v2.0.4-summer_edition

v2.0.4-summer_edition NOTES ↗ LATEST STABLE

EMBA v2.0.4 migrates to binwalk-ng, makes the local connector endpoint configurable, and expands CSV logging across all modules.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.0.4-summer_edition https://github.com/e-m-b-a/emba.git
# already have the repo? check out this version:
$ git checkout v2.0.4-summer_edition
```

Makes the local connector endpoint configurable via the change introduced in PR #2069.

Moves `MIN_TCP_SERV` to defaults and allows it to be adjusted via scanning profiles.

Migrates firmware extraction from binwalk to the actively maintained `binwalk-ng` fork.

Completes CSV output logging across all modules for structured, machine-readable results.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ

iOS Forensic Toolkit 10.11: bootloader extraction for Apple Watch Series 4 and 5, Apple T

BLOG ↗ RANK

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

> ElcomSoft's iOS Forensic Toolkit 10.11 extends bootloader-level full file system and keychain extraction beyond the A11 generation of Apple silicon, adding support for Apple Watch Series 4/5 and Apple TV 4K 2nd Gen, alongside updated OS compatibility.

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

01 **Bootloader-level extraction for Apple TV 4K 2nd Gen** NEW

78

Adds bootloader-level full file system and decrypted keychain extraction for Apple TV 4K 2nd Gen (A12 Bionic) via the `usbliter8` SecureROM exploit, recovering accounts, authentication tokens, paired device records, network history, and Siri-related artifacts. This is the first time iOS Forensic Toolkit's low-level extraction has been extended past the A11 generation of Apple silicon.

— Names exploit, chip generation, and specific recovered artifacts.

launch-20260921-e3070207 iOS Forensic Toolkit 10.11: bootloader …

THINNER COVERAGE BELOW

02 **Bootloader-level extraction for Apple Watch Series 4 and 5**

NEW

45

Adds bootloader-level extraction for Apple Watch Series 4 and Series 5, producing a full file system image and decrypted keychain.

— Names devices and output, no extraction mechanism given.

iOS Forensic Toolkit 10.11: bootloader …

03 **tvOS and audioOS version support updates** IMPROVED

30

Adds support for `tvOS 26.5` and `26.6`, and `audioOS 26.5` and `26.6`.

— Lists OS versions supported but no functional detail. launch-20260921-e3070207

> iOS Forensic Toolkit 10.11: bootloader extraction for Apple Watch Series 4 and 5, Appl

BLOG ↗ LATEST STABLE

iOS Forensic Toolkit 10.11 brings bootloader-level full file system and keychain extraction to Apple Watch Series 4/5 and Apple TV 4K 2nd Gen.

Adds bootloader-level extraction for Apple Watch Series 4 and Series 5, producing a full file system image and decrypted keychain.

Adds bootloader-level extraction for Apple TV 4K 2nd Gen, producing a full file system image and decrypted keychain.

Extends low-level extraction support past the A11 generation of Apple silicon for the first time.

announced · Sep 21, 2026 BLOG ↗

ElcomSoft: Adds bootloader-level keychain and full file system extraction for Apple TV 4K 2nd Gen (A12 Bionic) via the `usbliter8` SecureROM exploit, recovering accounts, authentication tokens, paired device records, network history, and Siri-related artifacts.

Adds support for tvOS 26.5 and 26.6.

Adds support for audioOS 26.5 and 26.6.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

NETRESEC

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

Network forensics, packet sniffers and IT security products. Download NetworkMiner and other free software for network security analysis.

> Netresec's FlowCarp gained statistical C2 protocol classification with custom model training, plus a free public demo API classifying traffic against 600+ known protocols.

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

01 **Free demo API for C2 protocol identification** NEW 75

A free demo server at `demo.flowcarp.com` supports identifying 600+ known C2 protocols without a local FlowCarp install; a PCAP can be submitted directly via `curl --data-binary @dump.pcap https://demo.flowcarp.com`.

Submit a PCAP to the free demo server during triage to instantly identify C2 protocols across 600+ known families without a local FlowCarp install.

```
$ curl --data-binary @dump.pcap https://demo.flowcarp.com | jq
```

— Names the endpoint and protocol count but not classification mechanism.

changelog-20260922-5cb5070d

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

FlowCarp adds statistical C2 protocol classification and a free demo API at `demo.flowcarp.com` supporting 600+ protocols

└──▷ **TRY IT**

Train a custom protocol model from sandbox C2 captures so FlowCarp can recognize a previously unknown malware family in future PCAPs.

```
$ flowcarp --input xyz.pcap --output
model:MalwareXYZ:MALICIOUS:malwarexyz.fcp
```

Classify traffic in a new PCAP against your custom model to confirm C2 protocol matches across different samples or sandbox runs.

```
$ flowcarp --model malwarexyz.fcp --input new.pcap --output -
```

Submit a PCAP to the free demo server during triage to instantly identify C2 protocols across 600+ known families without a local FlowCarp install.

```
$ curl --data-binary @dump.pcap https://demo.flowcarp.com | jq
```

WAS THIS USEFUL? 👍 👎

tcpdump

Feedback ⓘ tcpdump-4.99.7 CODE ↗ RANK

the TCPdump network dissector

> tcpdump 4.99.7 deepens MPTCP option dissection, adds TCP AccECN flag decoding, and warns when `--count` is misused on the command line.

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

01 **Expanded MPTCP option dissection** IMPROVED 64

Adds dissection of the MPTCP `MP_TCPRST` sub-option (subtype `0x8`), including the `U` / `V` / `W` / `T` flags and reset reasons such as Middlebox interference. Also parses the `data_len` field in the MPTCP `MP_CAPABLE` option and prints all eight capability flags (`A` through `H`).

— Names sub-option, subtype code, and all flags; no usage example given tcpdump-4.99.7

THINNER COVERAGE BELOW

02 **Warning for misused --count flag** IMPROVED 51

Adds an error message when `--count` is used without `-r` or `-V`, flagging a common command-line mistake.

— Names exact flags involved; clear condition to test tcpdump-4.99.7

03 **TCP AccECN flag decoding** NEW 34

Adds support for decoding the TCP `AE` (AccECN) flag in TCP header output.

— Names the flag but gives no mechanism or example output tcpdump-4.99.7

tcpdump-4.99.7 CODE ↗ LATEST STABLE

tcpdump 4.99.7 extends MPTCP dissection, adds TCP AccECN flag support, and warns on misused `--count`.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch tcpdump-4.99.7 https://github.com/the-tcpdump-group/tcpdump.git
# already have the repo? check out this version:
$ git checkout tcpdump-4.99.7
```

Adds an error message when `--count` is used without `-r` or `-V`.

> Adds dissection of the MPTCP `MP_TCPRST` sub-option (subtype `0x8`), including `U` / `V` / `W` / `T` flags and reset reasons such as `Middlebox interference`.
>
> Parses the `data_len` field in the MPTCP `MP_CAPABLE` option and prints all eight capability flags (`A` through `H`).
>
> Adds support for decoding the TCP `AE` (AccECN) flag in TCP header output.

WAS THIS USEFUL?  

SO-CRATES

Feedback ⓘ v4.2.0 NOTES ↗ RANK

SO-CRATES: Security Onion Containerized Rapid Analysis of Threats, Evil, and Sus!

> SO-CRATES v4.2.0 focuses on hardening deployment and API security, adding a reverse-proxy hostname allowlist, enforcing JSON content-type on API calls, and restricting default port exposure.

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

01 ALLOWED_HOSTS environment variable for reverse proxies

78

BREAKING

Adds the `ALLOWED_HOSTS` environment variable (e.g. `-e ALLOWED_HOSTS='*.killercoda.com'`) to support serving SO-CRATES behind a reverse proxy or cloud hostname. Browsing via IP or localhost requires no change, but deployments behind a proxy or cloud hostname must now set this variable or they will error out.

Run SO-CRATES behind a Killercoda (or other cloud/reverse-proxy) hostname so the host validation passes.

```
$ docker run -e ALLOWED_HOSTS='*.killercoda.com' so-crates
```

— Names exact env var and example command, but no deeper mechanism given v4.2.0

02 Localhost-only port binding in compose files

BREAKING

65

Docker/Podman compose files now publish the port on localhost only by default; to allow other hosts on the network to reach SO-CRATES, change the port mapping to `"8000:8000"`.

— Exact config change given, but limited explanation of underlying mechanism v4.2.0

THINNER COVERAGE BELOW

03 JSON API content-type enforcement

BREAKING

45

Scripts calling the JSON API must now send `Content-Type: application/json`; requests omitting this header will fail where they previously succeeded.

— Clear requirement stated but no example request or endpoint given v4.2.0

└──▷ **BREAKING ON UPGRADE**

! Serving SO-CRATES behind a reverse proxy or cloud hostname now requires setting the `ALLOWED_HOSTS` environment variable (e.g. `-e ALLOWED_HOSTS='*.killercoda.com'`); omitting it will produce an error.

! Docker/Podman compose files now bind the port to localhost only; setups relying on other machines reaching SO-CRATES on the network must change the port mapping to `"8000:8000"`.

! Scripts calling the JSON API must now send `Content-Type: application/json`; requests without this header will break.

v4.2.0 NOTES ↗ LATEST STABLE

SO-CRATES v4.2.0 adds `ALLOWED_HOSTS` env var for reverse-proxy deployments and tightens JSON API content-type enforcement.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v4.2.0 https://github.com/dougburks/so-crates.git
# already have the repo? check out this version:
$ git checkout v4.2.0
```

└──▷ **TRY IT**

Run SO-CRATES behind a Killercoda (or other cloud/reverse-proxy) hostname so the host validation passes.

```
$ docker run -e ALLOWED_HOSTS='*.killercoda.com' so-crates
```

Adds `ALLOWED_HOSTS` environment variable (e.g. `-e ALLOWED_HOSTS='*.killercoda.com'`) to support serving SO-CRATES behind a reverse proxy or cloud hostname; browsing via IP or localhost requires no change.

Requires scripts calling the JSON API to send `Content-Type: application/json`; requests omitting this header will now fail.

Docker/Podman compose files now publish the port on localhost only by default; change the port mapping to `"8000:8000"` to allow other hosts on the network.

└──▷ **BREAKING ON UPGRADE**

! Serving SO-CRATES behind a reverse proxy or cloud hostname now requires setting the `ALLOWED_HOSTS` environment variable (e.g. `-e ALLOWED_HOSTS='*.killercoda.com'`); omitting it will produce an error.

! Docker/Podman compose files now bind the port to localhost only; setups relying on other machines reaching SO-CRATES on the network must change the port

mapping to `"8000:8000"`.

! Scripts calling the JSON API must now send `Content-Type: application/json`; requests without this header will break.

WAS THIS USEFUL? 👍 👎

◆ **Firewalls & Network Defense**

Sophos Switch

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗

RANK

Sophos Switch is a network orchestration platform enabling centralized management and security policy enforcement across switching infrastructure.

> Sophos Switch now surfaces switch names as the primary identifier across management pages and lets admins search by name.

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

01 **Switch names as primary identifier with search** IMPROVED

50

Switch, site, and stack management pages in Sophos Central now display switch names as the primary identifier instead of other identifiers, and the search box on these pages supports searching for switches by name.

— Names the exact pages affected but no deeper mechanism given

changelog-20260922-7c747c22

> changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE
>
> Sophos Switch now displays switch names as the primary identifier on management pages, with name-based search.
>
> Displays switch names as the primary identifier on the switch, site, and stack management pages in Sophos Central.
>
> Enables searching for switches by name in the search box on switch, site, and stack management pages.

WAS THIS USEFUL? 👍 👎

Aviatrix

Feedback ⓘ announced · Sep 22, 2026 PRESS ↗ RANK

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

> Aviatrix launched Harvest and Decrypt Protection, a single software policy combining post-quantum ML-KEM encryption, Communication Governance, and a two-level Containment Assessment tool, alongside a new crypto visibility inventory and a free tier for early adopters.

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

01 **Post-quantum ML-KEM encryption with crypto agility** NEW

70

Aviatrix adds post-quantum key exchange using the `ML-KEM` standard on the control plane across Aviatrix-encrypted paths, making intercepted traffic unreadable to future quantum computers without a hardware refresh or network redesign. The post-quantum algorithm is exposed as a configuration setting — crypto agility — so it can be swapped via a configuration push rather than a hardware refresh cycle.

— Names the ML-KEM standard and swap mechanism, but no exact config command

launch-20260922-5ef39c4a

02 **Containment Assessment tool with level 2 runtime analysis**

NEW

70

A self-service Containment Assessment tool tests five properties of the cloud estate and returns a provisional Blast Radius and harvest exposure read in approximately five minutes. Containment Assessment level 2 extends this with a read-only, agentless engagement against live runtime flows that quantifies exposure as Reachable Value at Risk in dollars, before any policy is written.

— Concrete metrics and timing given, but no exact access path launch-20260922-5ef39c4a

THINNER COVERAGE BELOW

03 **Communication Governance workload policy enforcement** NEW

55

Communication Governance runs on the same gateways that handle encryption, enforcing explicit workload-to-workload reachability policy at the granularity of workload identity and protocol on every available path.

— Describes policy granularity but no config or API surface given

launch-20260922-5ef39c4a

04 **Free tier and trial for Harvest and Decrypt Protection** NEW

45

Harvest and Decrypt Protection is available at no cost for the first five policies and five nodes for existing customers, with no license and no time limit; new customers receive a 30-day free trial.

— Concrete numeric limits but no setup steps launch-20260922-5ef39c4a

05 **Cryptographic traffic inventory (crypto visibility)** NEW

43

Extends the quantum-safe roadmap with crypto visibility: a network-path cryptographic inventory mapping which applications communicate in cleartext, what they depend on, and where they communicate.

— High-level description with no interface or command detail launch-20260922-5ef39c4a

> announced · Sep 22, 2026 PRESS ↗
>
> Aviatrix launches Harvest and Decrypt Protection: post-quantum ML-KEM encryption plus Communication Governance in one software policy, free for first five policies.
>
> Adds post-quantum key exchange using the `ML-KEM` standard on the control plane across Aviatrix-encrypted paths, making intercepted traffic unreadable to future quantum computers without a hardware refresh or network redesign.
>
> Introduces Communication Governance on the same gateways that handle encryption, enforcing explicit workload-to-workload reachability policy at the granularity of workload identity and protocol on every available path.
>
> Provides crypto agility as a configuration setting — the post-quantum algorithm can be swapped with a configuration push rather than a hardware refresh cycle.
>
> Adds a self-service Containment Assessment tool that tests five properties of the cloud estate and returns a provisional Blast Radius and harvest exposure read in approximately five minutes.
>
> Adds a Containment Assessment level 2 — a read-only, agentless engagement against live runtime flows that quantifies exposure as Reachable Value at Risk in dollars, before any policy is written.
>
> +2 more

WAS THIS USEFUL? 👍 👎

Palo Alto Networks

Feedback ⓘ changelog · Sep 22, 2026

PRESS ↗ BLOG ↗ RANK

Palo Alto Networks provides cybersecurity platforms for threat prevention, detection, and response across networks, cloud, and endpoints.

> Palo Alto Networks launched Unit 42 Continuous Frontier AI Defense, an always-on agentic offensive security service that combines a continuous testing engine, a multi-model AI harness, and adversary simulation to find, validate, and remediate enterprise exposures.

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

01 **Continuous Testing Engine** NEW 65

Performs a full-estate baseline scan followed by always-on testing as an organization's environment changes, covering web apps, APIs, cloud infrastructure, source code repositories, and network assets.

— Clear mechanism and scope, no runnable action given 2026-09-22-d8eee514

02 **Multi-Model AI Harness for offensive security tasks** NEW 60

Routes offensive security tasks to the model best suited for each job, addressing the finding that no single model catches more than 40% of vulnerabilities and that leading models share less than 10% overlap in identified exposures.

— Mechanism and stats given but no way to act on it 2026-09-22-d8eee514

THINNER COVERAGE BELOW

03 **Unit 42 Continuous Frontier AI Defense service** NEW 58

A new always-on agentic offensive security service that continuously finds, validates, and remediates enterprise exposures using proprietary multi-model AI harnesses powered by Anthropic Claude Mythos 5 and OpenAI GPT-5.6-Cyber.

— Names the models and purpose but no usage path 2026-09-22-d8eee514

04 **Advanced Adversary Simulation** NEW 50

Proves real-world exploitability by validating end-to-end attack paths across first- and third-party web apps, APIs, cloud infrastructure, source code repositories, and network assets.

— Describes purpose and scope but no concrete usage detail 2026-09-22-d8eee514

changelog · Sep 22, 2026 PRESS ↗ BLOG ↗

Palo Alto Networks launches Unit 42 Continuous Frontier AI Defense, an always-on agentic offensive security service powered by Anthropic Claude Mythos 5 and OpenAI GPT-5.6-Cyber.

Introduces Unit 42 Continuous Frontier AI Defense, a new always-on agentic offensive security service that continuously finds, validates, and remediates enterprise exposures using proprietary multi-model AI harnesses with Anthropic Claude Mythos 5 and OpenAI GPT-5.6-Cyber.

Adds a Continuous Testing Engine that performs a full-estate baseline scan followed by always-on testing as an organization's environment changes, covering web apps, APIs, cloud infrastructure, source code repositories, and network assets.

Introduces a proprietary Multi-Model AI Harness that routes offensive security tasks to the model best suited for each job, addressing the finding that no single model catches more than 40% of vulnerabilities and that leading models share less than 10% overlap in identified exposures.

Delivers Advanced Adversary Simulation that proves real-world exploitability by validating end-to-end attack paths across first- and third-party web apps, APIs, cloud infrastructure, source code repositories, and network assets.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Falco

Feedback ⓘ announced · Sep 22, 2026 BLOG ↗ RANK

Cloud Native Runtime Security

> Falco 0.45.0 introduces raw byte matching in rule conditions and output, reload control endpoints, and GPU resource hijacking detection rules, alongside a set of breaking changes covering driver API compatibility, modern eBPF memory usage, container configuration keys, and package upgrade behavior.

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

01 **Reload control endpoints** NEW 90

Adds a `GET /reload` webserver endpoint reporting process identity, reload generation, and readiness so integrations can verify a completed reload, plus a new `reload_control` configuration key that opens a Unix socket accepting `POST /reload` for local reload requests (disabled by default).

— Names exact endpoint, method and config key launch-20260922-8850a46f

02 **Container plugin engine_timeout option** NEW 90

Adds an `engine_timeout` option to the container plugin (defaults to 10 seconds) to prevent an unresponsive runtime socket from blocking Falco startup indefinitely; set to `0` to disable.

— Exact flag, default value and disable option given launch-20260922-8850a46f

03 **Raw byte matching in rule conditions and output** BREAKING 85

Adds `\xHH` byte escape syntax in rule conditions for matching filenames or process arguments containing arbitrary bytes not representable as valid UTF-8, and now preserves raw bytes when extracting field values for rule matching rather than sanitized replacements. Both text and JSON alert output now escape control characters and replace invalid UTF-8 sequences for printable output; rule conditions and output consumers that relied on the old UTF-8 sanitization behavior will behave differently and should be reviewed after upgrading.

Write a rule that detects a process argument containing a specific non-UTF-8 byte sequence, now possible with raw byte escapes.

```yaml
- rule: Suspicious Raw Byte Argument
  desc: Detects a process argument containing a specific non-UTF-
8 byte
  condition: proc.args contains \xc0\xae
  output: 'Suspicious raw byte in argument (proc=%proc.name
args=%proc.args)'
  priority: WARNING
```

— Exact syntax plus runnable example rule given `launch-20260922-8850a46f`

04 **GPU resource hijacking detection rules** NEW 80

Adds two sandbox rules — Container Accessing GPU Device and GPU Management Tool Run in Container — for detecting GPU resource hijacking inside containers; both are disabled by default and require tuning the `user_known_gpu_workloads` list before enabling.

— Names both rules and the config list to tune `launch-20260922-8850a46f`

05 **Driver selection controls for package upgrades** BREAKING 80

Supports `FALCO_DRIVER_CHOICE=none` to skip driver setup and `FALCOCTL_DRIVER_VERSION` to specify an explicit driver version during DEB/RPM package upgrade transactions. DEB/RPM upgrades now also start the selected Falco service and install the driver automatically, so unattended upgrade scripts and custom driver pins must be reviewed to avoid unintended changes.

— Names exact env vars and the changed default behavior `launch-20260922-8850a46f`

06 **Modern eBPF driver buffer and drop-tracking improvements** 75

BREAKING

The modern eBPF probe now uses two auxiliary buffers per CPU to protect event construction on preemptible kernels, with graceful fallback and a warning on kernels lacking required BPF atomic operations — this doubles auxiliary map memory usage per CPU, so capacity should be reviewed before upgrading. New eBPF counters expose drops caused by reentrancy and exhausted auxiliary buffers, and CPU migrations, to help distinguish preemption-related capture pressure from other drop causes.

— Explains mechanism and memory impact but no direct action `launch-20260922-8850a46f`

07 **container_engines.* configuration removed** BREAKING 70

The `container_engines.*` configuration keys have been removed from the configuration schema; migrate to the container plugin configuration (e.g. `engines.cri.sockets`) instead.

— Names removed keys and the replacement config key launch-20260922-8850a46f

08 Driver API 11.0.0 compatibility enforcement

BREAKING

60

Falco 0.45.0 refuses to load kernel modules or probes not compatible with driver API 11.0.0 — updated drivers must be deployed before upgrading.

— Clear pre-upgrade action but limited mechanism detail launch-20260922-8850a46f

THINNER COVERAGE BELOW

09 Container plugin background metadata retrieval

IMPROVED

50

Container plugin 0.7.4 keeps already-collected container metadata and retrieves missing container info in the background when an initial listing is interrupted at startup.

— Describes behavior but no configuration surface named launch-20260922-8850a46f

└──▷ BREAKING ON UPGRADE

- ! Falco 0.45.0 refuses to load kernel modules or probes not compatible with driver API 11.0.0 — deploy updated drivers before upgrading.
- ! Modern eBPF now uses two auxiliary buffers per CPU, doubling auxiliary map memory per CPU — account for this before upgrading.
- ! Rule conditions that relied on UTF-8 sanitization before comparisons will behave differently; review and update custom rules accordingly.
- ! Output consumers that parsed Falco's formatted text or JSON output must account for the new control-character escaping and invalid-UTF-8 replacement.
- ! The `container_engines.*` configuration keys have been removed from the configuration schema — migrate to the container plugin configuration (e.g. `engines.cri.sockets`) instead.
- ! DEB/RPM upgrades now start the selected Falco service and install the driver — unattended upgrade scripts and custom driver pins must be reviewed to avoid unintended changes.

announced · Sep 22, 2026 BLOG ↗

Falco 0.45.0 adds raw byte matching in rules, reload control endpoints, and GPU detection rules

└──▷ USE IT

Write a rule that detects a process argument containing a specific non-UTF-8 byte sequence, now possible with raw byte escapes.

/etc/falco/rules.d/raw_byte.yaml

```
- rule: Suspicious Raw Byte Argument
  desc: Detects a process argument containing a specific non-UTF-8
byte
  condition: proc.args contains \xc0\xae
  output: 'Suspicious raw byte in argument (proc=%proc.name
args=%proc.args)'
  priority: WARNING
```

Adds `GET /reload` webserver endpoint reporting process identity, reload generations, and readiness so integrations can verify a completed reload.

Adds `reload_control` configuration key that opens a Unix socket accepting `POST /reload` for local reload requests (disabled by default).

Adds `\xHH` byte escape syntax in rule conditions for matching filenames or process arguments containing arbitrary bytes not representable as valid UTF-8.

Adds `engine_timeout` option to the container plugin (defaults to 10 seconds) to prevent an unresponsive runtime socket from blocking Falco startup indefinitely; set to `0` to disable.

Adds two sandbox rules — Container Accessing GPU Device and GPU Management Tool Run in Container — for detecting GPU resource hijacking inside containers (both disabled by default; tune `user_known_gpu_workloads` before enabling).

+6 more

└──▷ **BREAKING ON UPGRADE**

- ! Falco 0.45.0 refuses to load kernel modules or probes not compatible with driver API 11.0.0 — deploy updated drivers before upgrading.
- ! Modern eBPF now uses two auxiliary buffers per CPU, doubling auxiliary map memory per CPU — account for this before upgrading.
- ! Rule conditions that relied on UTF-8 sanitization before comparisons will behave differently; review and update custom rules accordingly.
- ! Output consumers that parsed Falco's formatted text or JSON output must account for the new control-character escaping and invalid-UTF-8 replacement.
- ! The `container_engines.*` configuration keys have been removed from the configuration schema — migrate to the container plugin configuration (e.g. `engines.cri.sockets`) instead.
- ! DEB/RPM upgrades now start the selected Falco service and install the driver — unattended upgrade scripts and custom driver pins must be reviewed to avoid unintended changes.

WAS THIS USEFUL?  

Bubblewrap

Feedback ⓘ v0.13.0 NOTES ↗ CODE ↗ RANK

Low-level unprivileged sandboxing tool used by Flatpak and similar projects

> Bubblewrap v0.13.0 adds mount flag remounting via mount_setattr(), tightens path argument validation, and introduces an opt-in debug logging build option.

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

01 Debug logging build option NEW

80

A new `-Ddebug_logging=true` Meson build option enables verbose debug logging in bubblewrap, gated at runtime by the `DEBUG_INVOCATION` environment variable so builds can ship with the capability without enabling it by default.

Build bubblewrap with debug logging enabled, then trigger verbose output for a sandboxed invocation to trace mount_setattr calls.

```
$ meson setup _builddir -Ddebug_logging=true
meson compile -C _builddir
DEBUG_INVOCATION=1 _builddir/bwrap --ro-bind /usr /usr --proc
/proc --dev /dev --unshare-pid bash
```

— Names exact build flag and env var with runnable example v0.13.0

THINNER COVERAGE BELOW

02 Empty path arguments now rejected BREAKING

55

Path arguments such as `--bind` and `--overlay-src` now reject empty strings instead of silently treating them as the root directory, changing prior behavior where an empty path implicitly meant `/`.

— Names affected flags and before/after behavior, no example v0.13.0

03 mount_setattr() for mount flag remounting NEW

25

Bubblewrap v0.13.0 adds use of mount_setattr() for remounting mount flags, per the release summary; no further mechanism or affected options are described.

— Only mentioned in summary, no detail or example given v0.13.0

v0.13.0 NOTES ↗ CODE ↗ LATEST STABLE

> Bubblewrap v0.13.0 adds mount_setattr() for mount flag remounting, rejects empty path args, and introduces `-Ddebug_logging=true` with `DEBUG_INVOCATION`

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.13.0
https://github.com/containers/bubblewrap.git
# already have the repo? check out this version:
$ git checkout v0.13.0
```

└──▷ **TRY IT**

Build bubblewrap with debug logging enabled, then trigger verbose output for a sandboxed invocation to trace mount_setattr calls.

```
$ meson setup _builddir -Ddebug_logging=true
meson compile -C _builddir
DEBUG_INVOCATION=1 _builddir/bwrap --ro-bind /usr /usr --proc /proc
--dev /dev --unshare-pid bash
```

Adds `-Ddebug_logging=true` build option to enable verbose debug logging, gated at runtime by the `DEBUG_INVOCATION` environment variable.

Path arguments such as `--bind` and `--overlay-src` now reject empty strings instead of silently treating them as the root directory.

WAS THIS USEFUL?  

◆ CNAPP

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Sep 22, 2026 | DOCS ↗ | BLOG ↗ | RANK

Rapid7 InsightCloudSec provides cloud security posture management and compliance monitoring across multi-cloud environments.

> Rapid7 InsightCloudSec's September release extends CrowdStrike Falcon Query Filters to Oracle Cloud Infrastructure and adds a wave of Azure detection capability spanning Cosmos DB authentication/encryption/network posture and MySQL/PostgreSQL Flexible Server private endpoint and public-access checks.

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

01 **Cosmos DB authentication, encryption and network posture detection**

NEW 90

Adds harvesting of the `disableLocalAuth` property and firewall IP rules (including `0.0.0.0` datacenter-access detection) for Azure Distributed Table (Cosmos DB) resources, plus three new Query Filters — Distributed Table Local Authentication Configuration (default: enabled), Distributed Table Encryption Configuration (default: provider-managed), and Distributed Table Without Network Traffic Restrictions — and three matching Insights: Distributed Table With Local Authentication Enabled, Distributed Table Encrypted With Provider Managed Key, and Distributed Table Without Network Traffic Restrictions.

— Names every harvested field, filter and insight with defaults 2026-09-22-687ac038

02 **Azure MySQL/PostgreSQL Flexible Server private endpoint and auth checks**

NEW 90

Adds harvesting of private endpoint connection counts for Azure Database for MySQL and PostgreSQL Flexible Servers, plus three new Query Filters — Database Instance Microsoft Entra-Only Authentication (default: not configured), Database Instance Flexible Private Endpoint Connections (configurable by engine type, default: MySQL Flexible Servers without private endpoints), and Database Instance Firewall Allows Public Ingress (flags a rule with start IP `0.0.0.0` and end IP other than `0.0.0.0`, excluding the Azure services-only `0.0.0.0`–`0.0.0.0` rule) — and five new Insights: Database Instance Without Microsoft Entra-Only Authentication (MySQL), Database Instance Without Private Endpoint (MySQL Flexible Server), `PostgreSQL Database Instance With Public Network Access Enabled`, `PostgreSQL Database Instance`

`Without Private Endpoint Connection`, and Database Instance Firewall Allows Public Ingress (SQL Server).

— Enumerates every filter and insight with concrete criteria `2026-09-22-687ac038`

03 **PostgreSQL log retention insight parameter matching** IMPROVED

75

Updates the Database Instance Log Retention Below Threshold (PostgreSQL) Insight to detect Azure PostgreSQL Flexible Servers where `logfiles.retention_days` is set to 3 days or fewer, and updates the Database Instance With Setting Below Threshold log_retention_days (PostgreSQL) Insight to match both the legacy `log_retention_days` and updated `logfiles.retention_days` parameter names using an OR condition.

— Names exact parameters and threshold value `2026-09-22-687ac038`

THINNER COVERAGE BELOW

04 **CrowdStrike Falcon Query Filter for OCI** NEW

53

Extends CrowdStrike Falcon agent Query Filters to Oracle Cloud Infrastructure (OCI) compute instances, joining existing AWS, Azure, and GCP coverage, enabling security teams to identify OCI instances by Falcon agent presence.

— Names platform coverage but not the filter's exact label `2026-09-22-687ac038`

changelog · Sep 22, 2026 DOCS ↗ BLOG ↗ LATEST STABLE

InsightCloudSec 26.9.22 adds CrowdStrike Falcon OCI support, seven new Azure Insights, and six new Query Filters for Cosmos DB, MySQL, and PostgreSQL.

Extends CrowdStrike Falcon agent Query Filters to Oracle Cloud Infrastructure (OCI) compute instances, joining existing AWS, Azure, and GCP coverage — enabling security teams to identify OCI instances by Falcon agent presence.

Adds harvesting of the `disableLocalAuth` property and firewall IP rules (including `0.0.0.0` datacenter-access detection) for Azure Distributed Table (Cosmos DB) resources.

Adds harvesting of private endpoint connection counts for Azure Database for MySQL and PostgreSQL Flexible Servers.

Adds Distributed Table Local Authentication Configuration Query Filter to identify Cosmos DB accounts by whether local (key-based) authentication is enabled or disabled (default: enabled).

Adds Distributed Table Encryption Configuration Query Filter to identify Cosmos DB accounts by customer-managed vs. provider-managed encryption keys (default: provider-managed).

+14 more

WAS THIS USEFUL?  

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.43.0 NOTES ↗ CODE ↗ RANK

Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment.

> Prowler 5.43.0 replaces the pilot FedRAMP 20x frameworks with two unified 2026 frameworks, adds seven new AWS Inspector/FIPS checks with supporting IAM permissions, and rounds out Private Cloud with self-registration control, external provider onboarding, AWS GovCloud failover, and invitation-management fixes.

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

01 **Seven new AWS Inspector and FIPS checks** NEW 90

Adds `inspector2_coverage_scan_status_active`, `inspector2_coverage_recently_scanned` (configurable via `inspector2_max_days_since_last_scan`, default 3), `inspector2_active_findings_no_known_exploited_vulnerabilities`, `inspector2_active_findings_kev_within_due_date`, `inspector2_active_findings_within_max_age` (configurable via `inspector2_active_finding_max_age_days`, default 192), `elbv2_listener_fips_tls_enabled`, and `transfer_server_fips_security_policy_enabled`. Also adds `inspector2:BatchGetFindingDetails` to the Prowler additions policy and CloudFormation scan role so the KEV checks can pull CISA Known Exploited Vulnerability data directly from Inspector without an external feed; a missing permission now reports `MANUAL` instead of a false `FAIL`.

— names every check, config key, default, and IAM permission 5.43.0

02 **FedRAMP 20x 2026 compliance frameworks** BREAKING 85

Adds `fedramp_20x_ksi_2026`, a universal compliance framework covering the 46 FedRAMP Key Security Indicators across AWS, Azure, GCP, Kubernetes, and M365 from a single definition, and `fedramp_20x_frr_class_c_2026`, covering 158 FedRAMP Class C provider rules across the same providers. These replace the pilot per-provider frameworks `fedramp_20x_ksi_low_aws`, `fedramp_20x_ksi_low_azure`, and `fedramp_20x_ksi_low_gcp`; any automation or stored results referencing the pilot IDs must be updated.

— names frameworks and counts; migration path stated but not exact steps 5.43.0

03 **Self-registration toggle for Private Cloud** NEW 75

Adds `UI_SELF_REGISTRATION_ENABLED` environment variable for Prowler Private Cloud deployments; set to `false` to make sign-up invitation-only, hiding the Sign up link and Create organization option.

— exact env var and behaviour, no rollout scope given 5.43.0

04 **AWS partition bootstrap failover for GovCloud** NEW 75

When `PROWLER_AWS_PARTITION` is set, bootstrap STS calls now try up to two additional regions in the partition on connection errors or timeouts, so GovCloud scans succeed on networks that route only to one regional endpoint.

— named env var and mechanism, no exact region list 5.43.0

05 **Public endpoint signing for report downloads** NEW 67

Adds `DJANGO_OUTPUT_S3_AWS_PUBLIC_ENDPOINT_URL` environment variable to sign report download URLs against a browser-reachable storage host, enabling downloads on deployments where object storage is only reachable inside the container network.

— named env var with clear use case, no example given 5.43.0

06 **Invitation lifecycle fixes** IMPROVED 65

Expired invitations are now re-invitable: a pending invitation past its expiry is marked expired and a new invitation is created for the same email rather than returning a generic error; `filter[state__in]` on the invitations endpoint no longer returns a server error.

— names endpoint filter and before/after behaviour 5.43.0

THINNER COVERAGE BELOW

07 **SMN topic subscription check for Huawei Cloud** NEW 45

Adds `smn_topic_subscriptions` check for Huawei Cloud, reporting SMN topics that have no subscriptions configured.

— single check named but described only briefly 5.43.0

08 **Registry marketplace and external provider onboarding** NEW 42

Adds registry marketplace and external provider onboarding for Private Cloud, with permission-based access, schema-driven credentials, connection checks, and scan launch.

— described only in prose, no named endpoints or config 5.43.0

└──▷ BREAKING ON UPGRADE

! The pilot framework IDs `fedramp_20x_ksi_low_aws`, `fedramp_20x_ksi_low_azure`, and `fedramp_20x_ksi_low_gcp` are replaced by `fedramp_20x_ksi_2026`; any automation or stored results referencing the pilot IDs must be updated.

5.43.0 NOTES ↗ CODE ↗ LATEST STABLE

Prowler 5.43.0 adds FedRAMP 20x 2026 frameworks, seven AWS Inspector/FIPS checks, partition failover, and a self-registration toggle

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 5.43.0 https://github.com/prowler-cloud/prowler.git
# already have the repo? check out this version:
$ git checkout 5.43.0
```

Adds `fedramp_20x_ksi_2026` universal compliance framework covering the 46 FedRAMP Key Security Indicators across AWS, Azure, GCP, Kubernetes, and M365 from a single definition, replacing the pilot per-provider frameworks.

Adds `fedramp_20x_frr_class_c_2026` universal compliance framework covering 158 FedRAMP Class C provider rules across AWS, Azure, GCP, Kubernetes, and M365.

Adds seven new AWS checks: `inspector2_coverage_scan_status_active`, `inspector2_coverage_recently_scanned` (configurable via `inspector2_max_days_since_last_scan`, default 3), `inspector2_active_findings_no_known_exploited_vulnerabilities`, `inspector2_active_findings_kev_within_due_date`, `inspector2_active_findings_within_max_age` (configurable via `inspector2_active_finding_max_age_days`, default 192), `elbv2_listener_fips_tls_enabled`, and `transfer_server_fips_security_policy_enabled`.

Adds `inspector2:BatchGetFindingDetails` to the Prowler additions policy and CloudFormation scan role so KEV checks can pull CISA Known Exploited Vulnerability data directly from Inspector without an external feed; missing permission reports `MANUAL` instead of a false `FAIL`.

Adds `UI_SELF_REGISTRATION_ENABLED` environment variable for Prowler Private Cloud deployments; set to `false` to make sign-up invitation-only, hiding the `Sign up` link and `Create organization` option.

+5 more

└──▷ BREAKING ON UPGRADE

! The pilot framework IDs `fedramp_20x_ksi_low_aws`, `fedramp_20x_ksi_low_azure`, and `fedramp_20x_ksi_low_gcp` are replaced by `fedramp_20x_ksi_2026`; any automation or stored results referencing the pilot IDs must be updated.

WAS THIS USEFUL? 👍 👎

Wiz

Feedback ⓘ changelog · Sep 22, 2026 BLOG ↗ RANK

Wiz is a cloud security platform that identifies and prioritizes vulnerabilities and risks across cloud infrastructure and applications.

> Wiz shipped a GA WIN partners endpoint in its MCP server, letting partner AI agents pull live tenant context on demand during investigations.

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

01 GA WIN partners endpoint for MCP agent integrations NEW

70

Wiz added a GA WIN (Wiz Integration Network) partners endpoint in the Wiz MCP, letting partner AI agents — Claude Code, Cursor, GitHub Copilot, or any remote-MCP-capable agent — search WIN documentation, retrieve API operations and required scopes, and query demo environment data without manual API calls. This introduces MCP-based agent integrations in WIN that let partner agents pull live Wiz context — Issues, Findings, affected resources, attack paths, and blast radius — directly from customer tenants on demand during investigations.

— Names the endpoint and agents but no method, path, or exact config. 2026-09-22-55728d69

> changelog · Sep 22, 2026 BLOG ↗
>
> Wiz MCP gains a GA WIN partners endpoint so partner AI agents can pull live Issues, Findings, and resources on demand.
>
> Adds a GA WIN partners endpoint in the Wiz MCP, letting partner AI agents (Claude Code, Cursor, GitHub Copilot, or any remote-MCP-capable agent) search WIN documentation, retrieve API operations and required scopes, and query demo environment data without manual API calls.
>
> Introduces MCP-based agent integrations in the Wiz Integration Network (WIN), enabling partner AI agents to pull live Wiz context — Issues, Findings, affected resources, attack paths, and blast radius — directly from customer tenants on demand during investigations.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Chainguard

Feedback ⓘ Week of 2026-09-21 NOTES ↗ RANK

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

> Chainguard launched a beta for storing and serving artifacts from an EU region, added eleven new container images across application, FIPS, base and commercial tiers, and flagged a breaking default-user change in the sonar-scanner-cli image that alters its working directory.

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

01 Default-user change breaks sonar-scanner-cli working directory

75

BREAKING

An October 12 default-user change in the `sonar-scanner-cli` image alters where the tool resolves its working directory, requiring pipelines to explicitly pass `-Dsonar.working.directory=/usr/src/.scannerwork -Dsonar.projectBaseDir=/usr/src` (e.g. `docker run --rm -v "$(pwd):/usr/src" cgr.dev/chainguard/sonar-scanner-cli -Dsonar.working.directory=/usr/src/.scannerwork -Dsonar.projectBaseDir=/usr/src`) to preserve existing behavior.

Preserve existing sonar-scanner-cli pipeline behavior after the October 12 default-user change by explicitly redirecting the working directory back to the project workspace.

```
$ docker run --rm -v "$(pwd):/usr/src" cgr.dev/chainguard/sonar-scanner-cli -Dsonar.working.directory=/usr/src/.scannerwork -Dsonar.projectBaseDir=/usr/src
```

— Gives exact runnable command and flags to work around the change

changelog-20260922-054520ad

02 **Sovereign Artifacts beta for EU data residency** NEW

65

Chainguard launched a Sovereign Artifacts beta that stores and serves Chainguard Containers and Chainguard Libraries artifact bytes in an EU region for enrolled customers, with automatic backfill of existing artifacts and approximately 30% faster pulls from Europe.

— Describes mechanism and benefit but no enrollment steps given Week of 2026-09-21

THINNER COVERAGE BELOW

03 **New container images across tiers** NEW

57

Eleven new container images were added: `cursor-agent` (application tier), `chainguard-desktop-workstation-pc` (base tier), `commercial-avocado` (commercial tier), `rustfs` (application tier), `cloudzero-agent` (application + FIPS tier), `trivy-k8s-wrapper` (application + FIPS tier), `access-log-exporter` (application + FIPS tier), `datahub-gms` (application + FIPS tier), `kates-connect` (application + FIPS tier), `kubeflow-trainer-controller-manager` (application + FIPS tier), and `tracee` (application + FIPS tier).

— Names every image and tier but no pull commands or usage detail Week of 2026-09-21

> Week of 2026-09-21 NOTES ↗ STABLE
>
> Chainguard adds EU Sovereign Artifacts beta for data residency and ships 18 new container images including FIPS variants.
>
> Launches Sovereign Artifacts beta, storing and serving Chainguard Containers and Chainguard Libraries artifact bytes in an EU region for enrolled customers, with automatic backfill of existing artifacts and approximately 30% faster pulls from Europe.
>
> Adds new `cursor-agent` (application tier) container image.
>
> Adds new `chainguard-desktop-workstation-pc` (base tier) container image.
>
> Adds new `commercial-avocado` (commercial tier) container image.
>
> Adds new `rustfs` (application tier) container image.
>
> +7 more

WAS THIS USEFUL? 👍 👎

Anchore

Feedback ⓘ

Extending Supply Chain Security to AI Models with Anchore Enterprise v6.2 (2026-09-15)

BLOG ↗ RANK

Anchore's software supply chain solutions automate vulnerability scanning, strengthen container security, and support compliance with NIST, FedRAMP & more.

> Anchore Enterprise v6.2 extends supply chain security into AI/ML models with native model artifact detection, an MCP Server for AI agent integration, VEX-aware policy suppression, and local SBOM generation during scanning.

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

01 **VEX-aware policy evaluation** NEW 70

Policies now support VEX annotations, automatically factoring in VEX status values such as `not_affected` or `fixed` to suppress false positives and prevent unnecessary build breaks.

— Names specific VEX status values and effect Extending Supply Chain Security to AI M…

02 **Anchore Enterprise MCP Server** NEW 65

A new lightweight container image, the Anchore Enterprise MCP Server, gives AI agents and tools direct, structured access to vulnerability findings, policy evaluation results, and SBOM catalog data without custom API integration scripts.

— Names the surface and data it exposes, no run steps

Extending Supply Chain Security to AI M…

03 **AI/ML model artifact detection in SBOM catalog** NEW 60

Anchore now natively inspects container images for AI/ML model artifacts, including `GGUF` files and Docker Model structures, indexing them as packages in the central SBOM catalog alongside traditional software components.

— Names formats and mechanism but no command shown

Extending Supply Chain Security to AI M…

THINNER COVERAGE BELOW

04 **Local SBOM generation during scanning** NEW 55

Anchore now supports local SBOM generation and output during image scanning workflows, letting developers inspect component inventories and compliance status before pushing artifacts to central registries.

— Describes workflow benefit but no exact command given

Extending Supply Chain Security to AI M…

05 **Severity-based package filtering** NEW 50

Packages within an application version can now be filtered by associated vulnerability severity (Critical, High, Medium) to accelerate targeted remediation.

— Names severity levels, clear UI filter starting point

Extending Supply Chain Security to AI M…

06 **Policy outcome change notifications** NEW 40

Automated notifications now trigger whenever a policy outcome changes at the application version level, alerting security teams when a previously compliant artifact becomes non-compliant.

— Describes trigger but no config or endpoint named

Extending Supply Chain Security to AI M…

Extending Supply Chain Security to AI Models with Anchore Enterprise v6.2 (2026-09-15)

BLOG ↗ LATEST STABLE

Anchore Enterprise v6.2 adds AI/ML model SBOM detection, MCP Server for agent integration, VEX-aware policies, and local SBOM generation.

Natively inspects container images for AI/ML model artifacts, including `GGUF` files and Docker Model structures, indexing them as packages in the central SBOM catalog alongside traditional software components.

Introduces the Anchore Enterprise MCP Server, a lightweight container image that gives AI agents and tools direct, structured access to vulnerability findings, policy evaluation results, and SBOM catalog data without custom API integration scripts.

Adds policy support for VEX annotations, allowing policies to automatically factor in VEX status values such as `not_affected` or `fixed` to suppress false positives and prevent unnecessary build breaks.

Triggers automated notifications whenever a policy outcome changes at the application version level, alerting security teams when a previously compliant artifact becomes non-compliant.

Enables filtering of packages within an application version by associated vulnerability severity (Critical, High, Medium) to accelerate targeted remediation.

+1 more

WAS THIS USEFUL? 👍 👎

◆ SIEM

Datadog Cloud SIEM

Feedback ⓘ changelog · Sep 22, 2026

DOCS ↗ CHANGELOG ↗ RANK

Datadog Cloud SIEM detects and investigates security threats across cloud infrastructure using logs, metrics, and threat intelligence.

> Datadog Cloud SIEM's biggest addition this window is Dynamic Severity rules with fine-grained access controls, alongside new suppression rules, below-threshold detection logic, and three new out-of-the-box integrations.

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

01 **Dynamic Severity rules for signals** NEW 75

Cloud SIEM adds Dynamic Severity rules that adjust the severity of security signals based on the assets they affect, with customizable severity levels, custom tags, and rule-scoped isolation. Access to these rules can be restricted with role-, team-, and user-level controls, requiring at least one editor role the current user belongs to before saving — done from the rule list's vertical three-dot menu via Restrict access.

Scope a Dynamic Severity rule's edit access to specific roles or teams so only authorized staff can modify signal severity adjustments.

> 📍 In Cloud SIEM, open the Dynamic Severity rule list, click the vertical three-dot menu for the rule, select Restrict access, and use the dropdown to add one or more roles, teams, or users — ensuring your own role is included before saving.

— Mechanism explained and UI steps given, no config keys 2026-09-22-9fd97243

02 **Below-threshold cases in Threshold rules** NEW 60

Threshold detection rules now support below-threshold detection cases (`a < N` and `a <= 0`) alongside the existing above-threshold case (`a > N`).

— Exact operators named but no rule config path shown 2026-09-22-9fd97243

THINNER COVERAGE BELOW

03 **New out-of-the-box integrations** NEW 40

Cloud SIEM adds out-of-the-box integrations for Palo Alto Networks Firewall, Windows Event Logs, and Wiz.

— Names integrations but no setup or config detail 2026-09-22-9fd97243

04 **Suppression rules for detections** NEW 30

Cloud SIEM adds new Suppression rules for Cloud SIEM detections.

— Named feature but no mechanism or usage given 2026-09-22-9fd97243

05 **Object Storage Request Estimation** NEW 20

Cloud SIEM adds an Object Storage Request Estimation capability.

— Bare mention, no mechanism or usage detail 2026-09-22-9fd97243

06 **Bits Investigation accuracy improvements** IMPROVED 10

Cloud SIEM improves Bits Investigation accuracy.

— Vague improvement claim with no specifics 2026-09-22-9fd97243

changelog · Sep 22, 2026 DOCS ↗ CHANGELOG ↗ LATEST STABLE

Datadog Cloud SIEM adds Dynamic Severity rules, Env Remapping, Code Channels in Slack, and Data Access Control for signals.

└──▷ **HOW TO FIND IT**

Scope a Dynamic Severity rule's edit access to specific roles or teams so only authorized staff can modify signal severity adjustments.

> 📍 In Cloud SIEM, open the Dynamic Severity rule list, click the vertical three-dot menu for the rule, select Restrict access, and use the dropdown to add one or more roles, teams, or users — ensuring your own role is included before saving.

Adds Dynamic Severity rules to adjust the severity of security signals based on the assets they affect, with customizable severity levels, custom tags, and rule-scoped isolation.

Adds role-, team-, and user-level access controls for Dynamic Severity rules, requiring at least one editor role the current user belongs to before saving.

Adds Object Storage Request Estimation capability.

Adds out-of-the-box integrations for Palo Alto Networks Firewall, Windows Event Logs, and Wiz.

Adds New Suppression rules for Cloud SIEM detections.

+2 more

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗

RANK

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

> Databricks shipped automatic sensitive-data classification in Unity Catalog, a GA account-wide query history system table, and a new Databricks Runtime 19 with SQL/DataFrame functions, session-scoped Python UDF env vars, and improved shuffle defaults, alongside expanded ABAC column mask coverage for VARIANT types.

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

01 Databricks Runtime 19 SQL/DataFrame functions and Python UDF env vars

NEW 93

Databricks Runtime 19 adds the `truncate` function to SQL, Scala DataFrame, and Python DataFrame APIs to truncate a number to a specified scale without rounding, and `trim_array` to remove elements from the end (positive count) or beginning (negative count) of an array. It also adds the `spark.pythonWorkerEnv.<variable-name>` configuration prefix to set session-scoped environment variables for Python UDF scalar worker processes, and enables auto-optimized shuffle v2 by default for workloads that do not set a fixed shuffle partition count.

Inject a custom environment variable (e.g. a secret path) into every Python UDF worker in a session without restarting the cluster.

```python
spark.conf.set("spark.pythonWorkerEnv.MY_SECRET_PATH",
"/dbfs/secrets/my_secret.txt")
```

Truncate a monetary amount to 2 decimal places without rounding, useful for audit-safe financial calculations.

```
$ SELECT truncate(price, 2) AS truncated_price FROM transactions;
```

Drop the last N elements from an array column to trim trailing nulls or sentinel values before aggregation.

```
$ SELECT trim_array(tags, 2) AS trimmed_tags FROM events;
```

— Runnable commands and exact config keys given `Databricks Runtime 19: September 21, 20…`

02 **Query history system table generally available** IMPROVED

70

The `system.query.history` system table is now GA, providing account-wide records of SQL warehouse and serverless compute executions, including execution status, duration metrics, and query source.

— Named table with scope but no query example `The query history system table is now g…`

03 **ABAC column masks extend VARIANT casts to ARRAY and MAP**

IMPROVED

60

VARIANT cast support in ABAC column masks now extends to `ARRAY` and `MAP` column types, in addition to previously supported `STRUCT` and other types, requiring Databricks Runtime 19 or above.

— Names types and requirement, no usage example `VARIANT cast extends to ARRAY and MAP c…`

THINNER COVERAGE BELOW

04 **Automatic sensitive data classification in Unity Catalog** NEW

35

New Data Classification feature automatically classifies and tags sensitive data in Unity Catalog.

— Names feature but no mechanism or config detail `changelog-20260922-22113423`

> `changelog · Sep 22, 2026` CHANGELOG ↗ LATEST STABLE
>
> Databricks adds automatic sensitive data classification and tagging in Unity Catalog.

> `Databricks Runtime 19: September 21, 2026 (2026-09-21)` BLOG ↗ STABLE
>
> Databricks Runtime 19 adds `truncate`, `trim_array` functions and session-scoped Python UDF worker env vars via `spark.pythonWorkerEnv`.
>
> └──▷ **USE IT**
>
> Inject a custom environment variable (e.g. a secret path) into every Python UDF worker in a session without restarting the cluster.

```python
spark.conf.set("spark.pythonWorkerEnv.MY_SECRET_PATH",
"/dbfs/secrets/my_secret.txt")
```

Truncate a monetary amount to 2 decimal places without rounding, useful for audit-safe financial calculations.

```
$ SELECT truncate(price, 2) AS truncated_price FROM transactions;
```

Drop the last N elements from an array column to trim trailing nulls or sentinel values before aggregation.

```
$ SELECT trim_array(tags, 2) AS trimmed_tags FROM events;
```

Adds `spark.pythonWorkerEnv.<variable-name>` configuration prefix to set session-scoped environment variables for Python UDF scalar worker processes.

Adds `truncate` function to SQL, Scala DataFrame, and Python DataFrame APIs to truncate a number to a specified scale without rounding.

Adds `trim_array` function to SQL, Scala DataFrame, and Python DataFrame APIs to remove elements from the end of an array (positive count) or the beginning (negative count).

Enables auto-optimized shuffle v2 by default for workloads that do not set a fixed shuffle partition count.

`The query history system table is now generally available (2026-09-21)` BLOG ↗

STABLE

The `system.query.history` system table is now generally available for account-wide query telemetry.

`VARIANT cast extends to ARRAY and MAP columns in ABAC column masks (2026-09-16)`

BLOG ↗ STABLE

VARIANT column masks in ABAC now cover ARRAY and MAP types alongside STRUCT, requiring Databricks Runtime 19+.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

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

> Tines 43.3.0 brings deeper case management into the storyboard, adding change tracking, formula-driven fields, and visual case creation.

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

01 **Case tracking, formulas, and visual creation in storyboard**

NEW 40

Tines now adds change control tracking for cases directly within the storyboard, formula support for case details fields, and the ability to create cases visually from within the storyboard.

.
.
.
.

— Names three case features but no config keys or exact steps.

changelog-20260922-6946ef4d

> changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE
>
> Tines 43.3.0 adds change control tracking and formula support for cases, plus visual case creation in the storyboard.
>
> Adds change control tracking for cases directly within the storyboard.
>
> Adds formula support for case details fields.
>
> Enables visual case creation from within the storyboard.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

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

> Huntress shipped Defender scan toggles for Managed Antivirus, macOS parity on the agent overview page, a new detection for a disclosed Screen Sharing exploit, and two Shadow Workflow hardening controls in Managed ISPM.

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

01 **Shadow Workflow hardening controls in Managed ISPM** NEW

83

Managed ISPM adds two new M365 security controls targeting Shadow Workflows: `Ensure Direct Send is disabled` for Microsoft Exchange Online, and `Tenant Allow/Block list should contain zero entries in the allow list` for Microsoft Defender for Office 365. Both controls can be scheduled for deployment via the Schedule for Deployment option rather than Managed Deployments to avoid unplanned end-user impact.

Schedule the Exchange Online Shadow Workflow hardening controls to roll out to tenants without auto-deploying to avoid unplanned end-user impact.

> 📍 In the Huntress console, open the Microsoft 365 hardening controls, locate `Ensure Direct Send is disabled` and 'Tenant Allow/Block list should contain zero entries in the allow list', and use the Schedule for Deployment option rather than Managed Deployments.

— Names both control titles and the exact deployment mechanism to use.

`New Security Controls in Managed ISPM` `changelog-20260922-a8927daf`

02 **Network and Archive File Scanning toggles for Defender** NEW

72

Managed Antivirus now exposes `Network File Scanning` and `Archive File Scanning` toggles for Microsoft Defender, found under EDR > Managed Antivirus > Defender Configuration > Scans.

Enable Archive File Scanning and explicitly disable Network File Scanning for Defender to harden malware-in-ZIP detection without impacting shared-drive performance.

> 📍 In the Huntress console, go to EDR › Managed Antivirus › Defender Configuration › Scans, set Network File Scanning to Off and Archive File Scanning to

On, then save.

— Exact UI path and toggle names given, no mechanism beyond that.

changelog-20260922-a8927daf

THINNER COVERAGE BELOW

03 **Detection for macOS Screen Sharing root exploit** NEW

50

A new detection fires when a Screen Sharing session attaches to a root user via a legacy authentication handshake, covering an unauthenticated file-read flaw in Apple's Screen Sharing service disclosed July 2026.

— Explains what triggers it but no user-facing action to take.

changelog-20260922-a8927daf

04 **macOS last logged-in user on agent overview** IMPROVED

43

The agent overview page now shows the last logged-in user for macOS endpoints, a field previously available only for Windows, requiring Agent v0.14.196 or later.

— Names version requirement but no navigation or command. changelog-20260922-a8927daf

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

Huntress adds macOS Screen Sharing visibility, quarantine-strip detections, and Defender scan toggle controls for M365 and endpoint hardening.

└──▷ **HOW TO FIND IT**

Enable Archive File Scanning and explicitly disable Network File Scanning for Defender to harden malware-in-ZIP detection without impacting shared-drive performance.

📍 In the Huntress console, go to EDR › Managed Antivirus › Defender Configuration › Scans, set Network File Scanning to Off and Archive File Scanning to On, then save.

Schedule the Exchange Online Shadow Workflow hardening controls to roll out to tenants without auto-deploying to avoid unplanned end-user impact.

📍 In the Huntress console, open the Microsoft 365 hardening controls, locate Ensure Direct Send is disabled and 'Tenant Allow/Block list should contain zero entries in the allow list', and use the Schedule for Deployment option rather than Managed Deployments.

Adds Network File Scanning and Archive File Scanning toggles for Microsoft Defender in Managed Antivirus, found under EDR > Managed Antivirus > Defender

Configuration > Scans.

Extends the agent overview page to show the last logged-in user for macOS endpoints (previously Windows-only), requiring Agent v0.14.196 or later.

Adds a detection that fires when a Screen Sharing session attaches to a root user via a legacy authentication handshake, covering an unauthenticated file-read flaw in Apple's Screen Sharing service disclosed July 2026.

New Security Controls in Managed ISPM NOTES ↗ STABLE

Huntress Managed ISPM adds two new M365 security controls targeting Shadow Workflows in Exchange Online and Defender for Office 365.

Adds `Ensure Direct Send is disabled` control for Microsoft Exchange Online, available to schedule for deployment in Managed ISPM.

Adds 'Tenant Allow/Block list should contain zero entries in the allow list' control for Microsoft Defender for Office 365, available to schedule for deployment in Managed ISPM.

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Service Gateway: Smart Protection Services version 2.0.11 (Sept 2026) (2026-09-22)

BLOG ↗ RANK

Trend Vision One is Trend Micro's cloud-based XDR platform providing unified security visibility and threat detection across endpoints, networks, and cloud environments.

> Trend Vision One's Service Gateway Smart Protection Services 2.0.11 release adds support for the TrendAI Vision One agent.

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

01 **TrendAI Vision One agent support on Service Gateway** NEW

30

Service Gateway Smart Protection Services version 2.0.11 adds support for the TrendAI Vision One™ agent.

— Names the agent and component but no mechanism or setup steps

Service Gateway: Smart Protection Servi…

> Service Gateway: Smart Protection Services version 2.0.11 (Sept 2026) (2026-09-22)
>
> BLOG ↗ LATEST STABLE
>
> Service Gateway Smart Protection Services 2.0.11 adds support for the TrendAI Vision One agent.

WAS THIS USEFUL? 👍 👎

Fleet

Feedback ⓘ fleet-v4.92.0 NOTES ↗ RANK

Open device management

> Fleet shipped a full Reports feature — scheduling, live queries, CSV export, label targeting and configurable result caps — alongside a broad 4.92.0 release that adds a cancelable Apple MDM command endpoint, per-platform disk encryption controls, Android and iOS MDM enhancements, and osquery config bandwidth-saving options, plus two breaking changes to log formatting and Windows profile resend behavior.

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

01 **Reports feature with scheduling, live queries and CSV export**

NEW 95

Fleet adds a Reports section that lets admins create, schedule, run live, and export reports as CSV. Storage is controlled via `server_settings.report_cap` (default 1,000 stored results per report) and `server_settings.discard_reports_data` globally, or per-report via the `discard_data` field on the create/update report API and per-report/global Store data toggles in Advanced options. `schedule_splay_percent` in agent options randomizes each host's check-in time (up to 10% by default) to spread load. Label-based host targeting (Fleet Premium) is available via a Custom target selector, a Manage automations bulk UI (admin-only) configures which reports send data to a log destination, and running `SELECT * FROM osquery_schedule` as a live query shows exactly when a host will next check in for a scheduled report.

Cap stored report results at 500 and disable global result persistence to protect database performance as you add more reports.

```json
{
  "server_settings": {
    "report_cap": 500,
    "discard_reports_data": true
  }
}
```

Spread scheduled report check-ins across a wider window to reduce simultaneous host load on a large fleet.

```yaml
schedule_splay_percent: 20
```

— Names every config key, field and query; runnable examples given `product docs`

02 **Osquery config bandwidth-saving options** NEW 95

Adds `osquery.config_in_memory_cache` to control an in-memory cache of the osquery config's scheduled-report section (disabled by default, set `true` to re-enable). Adds `osquery.config_etags` (env `FLEET_OSQUERY_CONFIG_ETAGS`, default off) to enable conditional request (etag) support on `/api/osquery/config`, returning a minimal `{"etag":"ok"}` response when config is unchanged. Adds `osquery.redis_config_etags` (env `FLEET_OSQUERY_REDIS_CONFIG_ETAGS`, default off, requires `osquery.config_etags`) to answer matching etag check-ins directly from a Redis-backed ETag store, skipping config build and database reads entirely.

Reduce agent config bandwidth on large fleets by enabling etag-based conditional config responses so unchanged configs return only a minimal acknowledgement.

```yaml
osquery:
  config_etags: true
  redis_config_etags: true
```

— Every config key and env var named with mechanism and example `fleet-v4.92.0`

03 **Cancelable pending Apple MDM commands** NEW 90

Adds `DELETE /api/v1/fleet/hosts/:id/commands/:command_uuid` to cancel a pending Apple MDM lock, wipe, clear passcode, or enable lost mode command before the host receives it.

Cancel a pending Apple MDM wipe before a host receives it — useful when a wipe was sent to the wrong device.

```
$ curl -X DELETE
https://fleet.example.com/api/v1/fleet/hosts/42/commands/a1b2c3d4
-e5f6-7890-abcd-ef1234567890 -H 'Authorization: Bearer <token>'
```

— Exact endpoint, method and runnable curl example provided `fleet-v4.92.0`

04 **Expanded Android host vitals collection** NEW 80

Adds collection of additional Android host vitals — USB debugging, passcode set, Google Play Protect, encryption status, manufacturer, security update version, kernel and bootloader version, software update status, API level, security posture, and per-SIM phone numbers — returned by get host endpoints.

— Long named list of vitals but endpoint path not given exactly `fleet-v4.92.0`

---

05 **SSO-required end user authentication for Fleet Desktop** NEW

70

Adds a new End user authentication setting (`fleet_desktop.sso_enabled` , Fleet Premium) requiring end users to sign in via SSO before accessing Fleet Desktop's My device page.

— Named config key and exact UI page it gates `fleet-v4.92.0`

---

06 **Per-platform disk encryption enforcement and key escrow** NEW

65

Adds per-platform disk encryption settings: enforcement and key escrow for macOS, enforcement for Windows, and key escrow for Linux.

— Scope per platform given but no config key or path named `fleet-v4.92.0`

---

07 **Inactive user status and activity timestamps** NEW 65

Adds an Inactive status with tooltip to the Users table, with new `last_login_at` , `last_activity_at` , and server-computed `status` fields (`active` , `inactive` , or `no_access`) returned from the users API.

— Named fields and enum values but no example request `fleet-v4.92.0`

---

08 **Human-readable duration fields in server logs** BREAKING

65

Duration fields in Fleet server logs (e.g. `took`) now render as human-readable strings with time units (e.g. `"1.116187ms"`) instead of raw nanosecond numbers — log pipelines that parse these fields numerically must be updated.

— Exact field, before/after format, and migration note given `fleet-v4.92.0`

---

09 **Windows user-scoped profiles not auto-resent after upgrade**

65

BREAKING

Windows user-scoped configuration profiles that already failed on an earlier Fleet version are not resent automatically after upgrading to 4.92.0 — admins must resend them once (or edit the profile) for Fleet to deliver them when a user is signed in.

— Clear before/after with an explicit remediation step `fleet-v4.92.0`

---

10 **MDM enrollment fields on host API responses** NEW 60

Adds `mdm.is_personal_enrollment` to host API responses, reporting whether the last MDM enrollment was personal (BYOD) and persisting after unenrollment. Adds `mdm.bootstrap_token_escrowed` to the get host and get host by Fleet Desktop token API responses.

— Both API fields named but no request example given `fleet-v4.92.0`

THINNER COVERAGE BELOW

11 **CVSS severity filter on vulnerability exposure dashboard** NEW 55

Adds a severity (CVSS score) filter to the Vulnerability exposure dashboard chart (Fleet Premium), with min/max score inputs appearing only when custom severity is selected.

— UI behaviour described but no exact navigation or API given `fleet-v4.92.0`

12 **Per-endpoint max request body size override** NEW 50

Adds `server.endpoint_request_size_overrides` configuration to set a max request body size per API endpoint.

— Config key named but no example or exact endpoints listed `fleet-v4.92.0`

13 **Experimental WebSocket notification transport** NEW 50

Adds `websocket.transport_enabled` server configuration to enable an experimental WebSocket notification transport for fleetd agents (disabled by default).

— Named config key but no usage detail beyond default state `fleet-v4.92.0`

14 **Automatic in-house app install on iOS/iPadOS enrollment** NEW 50

Adds automatic installation of in-house apps (`.ipa`) on iOS and iPadOS hosts when they enroll into Fleet.

— Describes behaviour with format but no setup steps `fleet-v4.92.0`

15 **Autopilot registry sync for pending Windows hosts** NEW 50

Adds Windows devices from a connected Microsoft Entra tenant's Autopilot registry to Fleet as pending hosts before enrollment.

— Mechanism named but no connection or config steps given `fleet-v4.92.0`

16 **Resend configuration profile on policy automation failure** NEW 40

Adds support for resending a configuration profile on policy failure as part of a policy automation.

– Bare description, no config key or trigger detail `fleet-v4.92.0`

17 **QR code for host enrollment link** NEW 40

Adds a QR code for the enrollment link on the Android and iOS/iPadOS tabs of the Add hosts modal.

– Exact UI location named though feature itself is minor `fleet-v4.92.0`

18 **display_name for Fleet-maintained apps in GitOps** NEW 35

Adds `display_name` support for Fleet-maintained apps in GitOps.

– Named field but no file path or usage shown `fleet-v4.92.0`

19 **Sort versions table by version on Software details page** NEW 30

Adds the ability to sort the versions table by version on the Software details page.

– UI path named but feature is a minor sort toggle `fleet-v4.92.0`

20 **Audit activity for unmatched SCIM deprovisioning** NEW 25

Adds an audit activity when a deprovisioned SCIM user has no email to match a Fleet account.

– Bare description with no field or endpoint named `fleet-v4.92.0`

└──▷ **BREAKING ON UPGRADE**

! Duration fields in Fleet server logs (e.g. `took`) now render as human-readable strings with time units (e.g. `"1.116187ms"`) instead of raw nanosecond numbers — log pipelines that parse these fields numerically must be updated.

! Windows user-scoped configuration profiles that already failed on an earlier Fleet version are not resent automatically after upgrading — admins must resend them once (or edit the profile) for Fleet to deliver them when a user is signed in.

`fleet-v4.92.0` NOTES ↗ STABLE

Fleet 4.92.0 adds Android MDM commands, per-platform disk encryption, cancelable Apple MDM commands, and new osquery bandwidth-saving config options.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch fleet-v4.92.0
https://github.com/fleetdm/fleet.git
# already have the repo? check out this version:
$ git checkout fleet-v4.92.0
```

└──▷ **USE IT**

Reduce agent config bandwidth on large fleets by enabling etag-based conditional config responses so unchanged configs return only a minimal acknowledgement.

```yaml
osquery:
  config_etags: true
  redis_config_etags: true
```

Cancel a pending Apple MDM wipe before a host receives it — useful when a wipe was sent to the wrong device.

```
$ curl -X DELETE
https://fleet.example.com/api/v1/fleet/hosts/42/commands/a1b2c3d4-
e5f6-7890-abcd-ef1234567890 -H 'Authorization: Bearer <token>'
```

Adds `osquery.config_in_memory_cache` server configuration option to control the in-memory cache of the osquery config's scheduled-report section (set to `true` to re-enable; disabled by default).

Adds `osquery.config_etags` server option (`FLEET_OSQUERY_CONFIG_ETAGS` , default off) to enable conditional request (etag) support on `/api/osquery/config` , sending a minimal `{"etag":"ok"}` response to agents whose configuration is unchanged.

Adds `osquery.redis_config_etags` server option (`FLEET_OSQUERY_REDIS_CONFIG_ETAGS` , default off, requires `osquery.config_etags`) to answer matching etag check-ins directly from a Redis-backed ETag store, skipping the config build and database reads entirely.

Adds `server.endpoint_request_size_overrides` configuration to set a max request body size per API endpoint.

Adds `websocket.transport_enabled` server configuration to enable an experimental WebSocket notification transport for fleetd agents (disabled by default).

+15 more

└──▷ **BREAKING ON UPGRADE**

- ! Duration fields in Fleet server logs (e.g. `took`) now render as human-readable strings with time units (e.g. `"1.116187ms"`) instead of raw nanosecond numbers — log pipelines that parse these fields numerically must be updated.
- ! Windows user-scoped configuration profiles that already failed on an earlier Fleet version are not resent automatically after upgrading — admins must resend them once (or edit the profile) for Fleet to deliver them when a user is signed in.

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ Recommended privacy permissions for iPhone and iPad

NOTES ↗ RANK

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

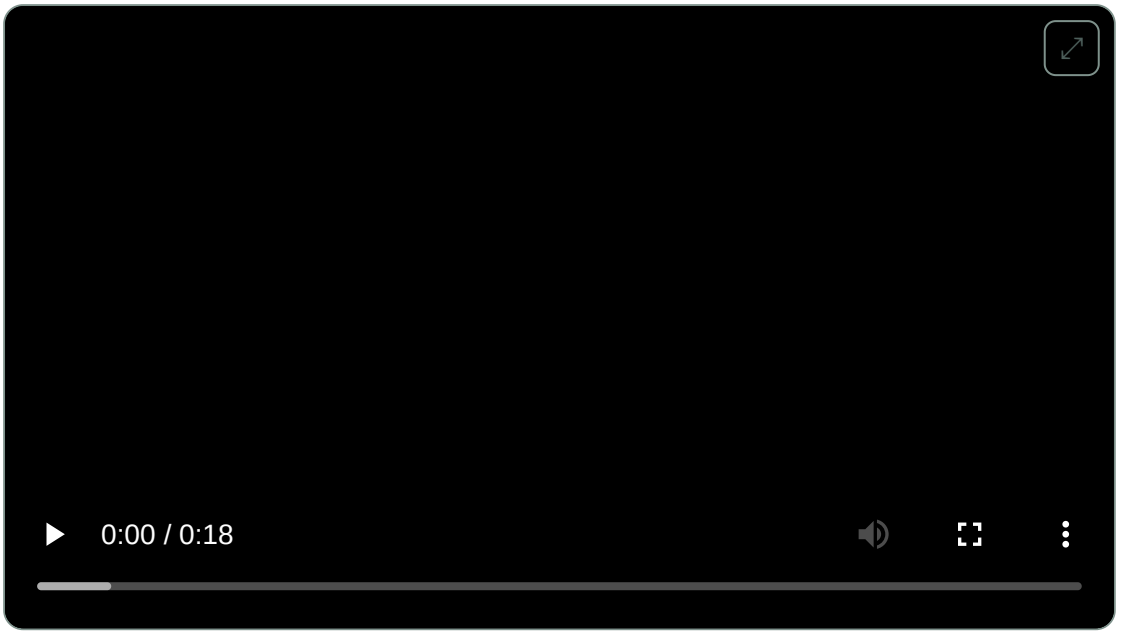
> Iru added a Recommended privacy permissions area to its Privacy Library Item for iPhone and iPad, letting admins batch-configure app and Safari permissions into a single consent prompt.

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

01 Recommended privacy permissions for iPhone and iPad NEW

58

A new Recommended privacy permissions area inside the Privacy Library Item lets admins specify app permissions and Safari website permissions that managed apps and sites require. End users see a single grouped consent prompt with Allow all or Not Now options instead of per-app popups.



When onboarding managed iPhones, configure recommended permissions so users receive a single grouped consent prompt instead of per-app popups.

> 📍 In the admin console, open the Privacy Library Item, navigate to the Recommended privacy permissions section for iPhone and iPad, add the required app permissions and Safari website permissions, then save and push to managed devices.

— Names UI area and behavior but no config keys or API.

Recommended privacy permissions for iPh…

Recommended privacy permissions for iPhone and iPad NOTES ↗ LATEST

STABLE

Iru's Privacy Library Item gains a Recommended privacy permissions area for iPhone and iPad, letting admins batch-prompt users for app and Safari permissions.

└──▷ HOW TO FIND IT

When onboarding managed iPhones, configure recommended permissions so users receive a single grouped consent prompt instead of per-app popups.

📍 In the admin console, open the Privacy Library Item, navigate to the Recommended privacy permissions section for iPhone and iPad, add the required app permissions and Safari website permissions, then save and push to managed devices.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Microsoft MSTICPy

Feedback ⓘ v3.1.0 NOTES ↗ RANK

Microsoft Threat Intelligence Security Tools

> MSTICPy 3.1.0 adds a new no-key threat-intelligence provider and Python 3.14 support, alongside two breaking changes to runtime dependencies and a user-agent header key.

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

01 **User-agent header key renamed** BREAKING 75

Code that directly indexes the dictionary returned by mp_ua_header() must replace `["UserAgent"]` with `["User-Agent"]`; code that passes the complete dictionary as HTTP headers does not need to change, and code needing the raw identifier can use `MSTICPY_USER_AGENT`.

– Exact old/new key names and env var given, actionable migration v3.1.0

02 **ScanMalware threat-intelligence provider** NEW 70

Adds the `ScanMalware` TI provider, supporting verdict lookups for hosts and URLs, passive-DNS context for IP addresses, and `quick` query modes for bulk lookups. No API credentials are required to use it.

– Names provider, lookup types, and quick mode but no usage example v3.1.0

THINNER COVERAGE BELOW

03 **NumPy 1.x dropped from runtime requirements** BREAKING 45

NumPy 1.x is no longer supported by the declared runtime requirements; users should review environment constraints for NumPy 2 compatibility before upgrading.

– Clear migration note but no version specifics v3.1.0

04 **Python 3.14 support** NEW 30

Adds Python 3.14 to the supported package classifiers and CI test matrix.

– Bare compatibility note, no further detail v3.1.0

└──▷ **BREAKING ON UPGRADE**

! NumPy 1.x is no longer supported by the declared runtime requirements — review environment constraints for NumPy 2 compatibility before upgrading.

! Code that directly indexes the dictionary returned by mp_ua_header() must replace `["UserAgent"]` with `["User-Agent"]`; code that passes the complete dictionary as HTTP headers does not need to change, and code needing the raw identifier can use `MSTICPY_USER_AGENT`.

v3.1.0 NOTES ↗ LATEST STABLE

MSTICPy 3.1.0 adds a new ScanMalware TI provider for DNS, IP, and URL observables with no API key required.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.1.0 https://github.com/microsoft/msticpy.git
# already have the repo? check out this version:
$ git checkout v3.1.0
```

Adds the `ScanMalware` threat-intelligence provider supporting verdict lookups for hosts and URLs, passive-DNS context for IP addresses, and `quick` query modes for bulk lookups — no API credentials required.

Adds Python 3.14 to the supported package classifiers and CI test matrix.

└──▷ BREAKING ON UPGRADE

! NumPy 1.x is no longer supported by the declared runtime requirements — review environment constraints for NumPy 2 compatibility before upgrading.

! Code that directly indexes the dictionary returned by mp_ua_header() must replace `["UserAgent"]` with `["User-Agent"]`; code that passes the complete dictionary as HTTP headers does not need to change, and code needing the raw identifier can use `MSTICPY_USER_AGENT`.

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 7.260921.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI's latest window adds new Prometheus telemetry metrics, granular health-monitoring and usage-metrics manager configuration, Redis TLS SNI support, and exponential back-off tuning for missing-reference retries, alongside product updates for Fintel reports, News Feed integrations, and user merge.

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

01 Configurable health-monitoring and usage-metrics manager NEW

95

Adds `app:health_monitoring:dependency_check_interval` (`APP__HEALTH_MONITORING__DEPENDENCY_CHECK_INTERVAL`) to control how often dependency health checks run, in milliseconds, for ElasticSearch/OpenSearch, S3/MinIO, RabbitMQ, and Redis, and `app:health_monitoring:usage_metrics_interval` (`APP__HEALTH_MONITORING__USAGE_METRICS_INTERVAL`) to control how often storage and ingestion usage metrics are collected and shared across the cluster via Redis, settable to `0` to disable collection. Adds `platform_usage_metrics_manager:enabled` (`PLATFORM_USAGE_METRICS_MANAGER__ENABLED`) to toggle the manager that gathers these usage metrics, and `platform_usage_metrics_manager:lock_key` (`PLATFORM_USAGE_METRICS_MANAGER__LOCK_KEY` , default `platform_usage_metrics_manager_lock`) to control its distributed lock.

Set health monitoring refresh intervals so Prometheus scrapes never trigger live dependency probes — tune dependency checks and expensive usage scans independently.

yaml

```yaml
app:
  health_monitoring:
    dependency_check_interval: 30
    usage_metrics_interval: 120
```

Configure health-monitoring check cadence and enable the usage-metrics manager in a multi-node OpenCTI deployment so each node refreshes dependency status every 10 s and collects storage/ingestion metrics every 60 s.

```yaml
app:
  health_monitoring:
    dependency_check_interval: 10000
    usage_metrics_interval: 60000
platform_usage_metrics_manager:
  enabled: true
  lock_key: platform_usage_metrics_manager_lock
```

— Exact config keys, env vars and defaults with runnable YAML examples. product docs

02 **Exponential back-off for missing-reference retries** IMPROVED 95

Adds `OPENCTI_MISSING_REF_RETRY_EXPONENTIAL` to switch missing-reference retries from a flat random 1-3s delay to an exponential back-off schedule (averages 0.5s, 1s, 2s, 4s with jitter) applied to workers and connectors via the `pycti` library, `OPENCTI_MISSING_REF_RETRY_INITIAL_DELAY` to set the average wait before the first retry (minimum 0), and `OPENCTI_MISSING_REF_RETRY_FACTOR` to set the back-off multiplier applied at each retry (minimum 1; 1 = constant wait).

Enable exponential back-off for missing-reference retries in a connector environment where referenced entities often arrive within a second, reducing unnecessary wait time on early retries.

```
$ OPENCTI_MISSING_REF_RETRY_EXPONENTIAL=true
OPENCTI_MISSING_REF_RETRY_INITIAL_DELAY=0.5
OPENCTI_MISSING_REF_RETRY_FACTOR=2
```

— Full env-var set with defaults and a runnable configuration example. product docs

03 **New health and usage Prometheus metrics** NEW 80

Adds `opencti_dependency_up` reporting connectivity state (up/down) for each dependency (`elasticsearch` , `storage` , `rabbitmq` , `redis`) as a Boolean, `opencti_elasticsearch_used_size_bytes` reporting total ElasticSearch/OpenSearch primary store size in bytes (replicas excluded), `opencti_storage_used_size_bytes` reporting total S3/MinIO bucket size in bytes, and `opencti_queue_consumers` reporting active consumers on push queues broken down by connector type (`EXTERNAL_IMPORT` , `INTERNAL_ENRICHMENT` , `INTERNAL_IMPORT_FILE` , `INTERNAL_EXPORT_FILE` , or `UNKNOWN`).

— Metrics named precisely but no scrape or query example given. product docs

04 **Redis TLS SNI server name config** NEW 70

Adds `redis:tls_servername` (`REDIS__TLS_SERVERNAME`) config key to supply an optional shared server name for TLS SNI and certificate validation when running Redis in cluster mode.

Add a Redis TLS server name for SNI validation when all cluster nodes share the same certificate CN, without per-node overrides.

```yaml
redis:
  tls_servername: redis.internal.example.com
```

— Named config key with a runnable example but thin mechanism detail. product docs

THINNER COVERAGE BELOW

05 **User merge reattributes relations and history** IMPROVED 45

Extends user merge to re-point assignee and participant relations to the target user, and to rewrite history and activity attribution to the target user.

— Describes the behaviour change but no command or UI path given. 7.260921.0

06 **News Feed support for OCTI integrations** NEW 40

Adds News Feed support for OCTI integrations (excluding connectors), surfacing a new integration resource type in the feed.

— Names the new resource type but gives no further mechanism. 7.260921.0

07 **Column selection in Fintel reports** NEW 35

Adds column selection for entity attributes in Fintel reports, letting analysts customize which fields appear in generated tables.

— Describes the capability but no UI path or config named. 7.260921.0

7.260921.0 NOTES ↗ LATEST STABLE

OpenCTI 7.260921.0 adds News Feed integration support, Fintel column selection, and user-merge history reattribution.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260921.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260921.0
```

> Adds column selection for entity attributes in Fintel reports, letting analysts customize which fields appear in generated tables.
>
> Adds News Feed support for OCTI integrations (excluding connectors), surfacing a new `integration` resource type in the feed.
>
> Extends user merge to re-point assignee and participant relations to the target user.
>
> Extends user merge to rewrite history and activity attribution to the target user.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Node v6.18.0: Parallel processing across transforms and sources (2026-09-22)

BLOG ↗ RANK

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

> Tenzir shipped a large batch of platform and node configuration surfaces this window — agent skills for AI coding tools spanning TQL and eight security schemas, four new TQL chart operators with dashboard export, an AWS CloudFormation deployment path for the Platform Sovereign Edition, centralized TLS/proxy/config-file controls, and pipeline parallelism and storage tuning knobs — alongside a v6.18.0 node release adding a recursive `search` function and removing the legacy non-neo executor.

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

01 **Agent skills for AI coding agents** NEW 98

Adds `npx skills add tenzir/skills` (also via `uvx skeel add` or `gh skill install`) to install Tenzir agent skills into Claude Code, Cursor, Codex, GitHub Copilot, and 40+ other coding agents from the `tenzir/skills` repository. The `tenzir` skill packages the full documentation (TQL, operators, functions, integrations, deployment, configuration) with progressive disclosure; `tenzir-ocsf` covers OCSF event classes, objects, attributes, profiles, and extensions; `tenzir-ecs` covers ECS fields, fieldsets, categorization values, custom fields, and OpenTelemetry alignment; `tenzir-cef` covers the ArcSight CEF standard including 174 predefined extension keys and 479 ESM data fields across 18 schema groups; `tenzir-leef` covers IBM QRadar LEEF 1.0/2.0 including 45 predefined event attributes and `devTime` / `devTimeFormat` patterns; `tenzir-cim` covers Splunk CIM 8.5 data models, datasets, fields, tags, constraints, and lookups; `tenzir-asim` covers Microsoft Sentinel ASIM schemas, fields, aliases, and roles; `tenzir-udm` covers Google SecOps UDM ingestion payloads (e.g. `metadata.eventType`) and YARA-L/Detect Engine logic (e.g. `$event.metadata.event_type`); and `tenzir-edm` covers the FortiSIEM 7.5.0 Event Data Model. Skills install to project scope by default (`./<agent>/skills/`) or globally with `-g` (`~/<agent>/skills/`), can target specific agents with `-a` (e.g. `-a claude-code -a cursor -a codex`), and are managed via `npx skills list`, `npx skills check`, `npx skills update`, and `npx skills remove`.

Install the OCSF schema skill globally so any project's coding agent can map raw events to OCSF event classes without re-installing per repo.

```
$ npx skills add -g tenzir/skills@tenzir-ocsf
```

Install the full Tenzir docs skill scoped to Claude Code and Cursor only, then keep it current — useful in a team repo where only those two agents are standardized.

```
$ npx skills add tenzir/skills@tenzir -a claude-code -a cursor
npx skills update
```

Audit which skills are installed globally and check whether any are out of date before updating a shared developer environment.

```
$ npx skills list -g
npx skills check
npx skills update
```

— Every skill, flag, and subcommand named with runnable commands `product docs`

02 **Chart operators for pipelines and dashboards** NEW 96

Adds four TQL chart operators: `chart_pie` (`label` / `x` , `value` / `y`), `chart_bar` (`x` , `y` , `group` , `position` e.g. `'stacked'` , `y_log`), `chart_line` (`x` , `y` , `group` , `resolution` , `x_min` , `x_max` , `y_min` , `y_max` , `fill`), and `chart_area` (same parameters as `chart_line` plus `position='stacked'`). All support inline aggregation expressions (e.g. y=mean(current_memory_usage)), removing the need for a preceding `summarize` ; `resolution` buckets time series with floor/ceiling alignment and implicit time-range filtering when combined with `x_min` / `x_max` ; `fill` replaces empty buckets (e.g. `fill=0`); and `group` splits data into series by field value without enumerating series names. Charts can be downloaded from the Explorer or a dashboard as PNG or SVG, and added to a dashboard via the Dashboard button in the Explorer for persistent, auto-refreshing visualization.

Visualize top application types from Zeek connection logs as a pie chart — useful for a quick protocol breakdown in the Explorer.

```
$ export
where @name == "zeek.conn"
where service != null
top service
head
chart_pie label=service value=count
```

Plot per-pipeline event throughput as a multi-series line chart over the last 24 hours, bucketed into 30-minute intervals — replaces a manual where + summarize chain.

```
$ metrics "publish"
chart_line (
  x=timestamp
  y=sum(events)
  x_min=now() - 1d
  group=pipeline_id
  resolution=30min
)
```

Show stacked bar chart of coin-flip outcomes across time-of-day groups to compare contributions while preserving cumulative totals.

```
$ from (
  {outcome: "tails", n: 30, time: "Morning"}
  {outcome: "heads", n: 20, time: "Morning"}
  {outcome: "tails", n: 25, time: "Afternoon"}
  {outcome: "heads", n: 35, time: "Afternoon"}
)
chart_bar x=outcome y=n group=time position="stacked"
```

— Every operator and parameter named with runnable pipeline examples `product docs`

03 **AWS CloudFormation deployment for Sovereign Edition** NEW

95

Deploys the Tenzir Platform Sovereign Edition on AWS via a publicly hosted CloudFormation template (`https://tenzir-marketplace-resources.s3.eu-west-1.amazonaws.com/tenzir-platform.yml`) that provisions VPC, ECS Fargate, Lambda, App Runner, RDS PostgreSQL, S3, Cognito, ACM, and Route53. External OIDC providers can replace Cognito via the `UseExternalOIDC`, `ExternalOIDCIssuerURL`, `ExternalOIDCClientID`, and `ExternalOIDCClientSecret` parameters, and `RandomSubdomain` auto-prefixes subdomains (e.g. `ui.abc123.example.org`) for

staging or multi-tenant setups. A `copy-images.sh` script pulls the three Marketplace images (`tenzir/se-platform-ui` , `tenzir/se-platform-api` , `tenzir/se-platform-gateway`) from the AWS Marketplace ECR registry (`709825985650.dkr.ecr.us-east-1.amazonaws.com`) into a local ECR account before stack creation.

Deploy the full Tenzir Platform stack on AWS in one command, pinning to a specific image version and using your own domain.

```
$ IMAGE_TAG=latest
AWS_REGION=$(aws configure get region)
AWS_ACCOUNT=$(aws sts get-caller-identity --query Account --
output text)
aws cloudformation create-stack \
  --stack-name tenzir-platform \
  --template-url https://tenzir-marketplace-resources.s3.eu-west-
1.amazonaws.com/tenzir-platform.yml \
  --parameters \
    ParameterKey=DomainName,ParameterValue=example.org \
    ParameterKey=RandomSubdomain,ParameterValue=false \
    ParameterKey=UseExternalOIDC,ParameterValue=false \

ParameterKey=ContainerImageUI,ParameterValue=$AWS_ACCOUNT.dkr.ecr
.$AWS_REGION.amazonaws.com/tenzir/se-platform-ui:$IMAGE_TAG \

ParameterKey=ContainerImageAPI,ParameterValue=$AWS_ACCOUNT.dkr.ec
r.$AWS_REGION.amazonaws.com/tenzir/se-platform-api:$IMAGE_TAG \

ParameterKey=ContainerImageGateway,ParameterValue=$AWS_ACCOUNT.dk
r.ecr.$AWS_REGION.amazonaws.com/tenzir/se-platform-
gateway:$IMAGE_TAG \
  --capabilities CAPABILITY_NAMED_IAM
```

Copy Marketplace container images to your own ECR before deploying — required before running the CloudFormation stack.

```
$ VERSION=v1.0.0 ./copy-images.sh
```

Integrate an external OIDC provider (e.g., Okta or Keycloak) instead of Cognito when deploying the stack.

```
$ aws cloudformation create-stack \
  --stack-name tenzir-platform \
  --template-url https://tenzir-marketplace-resources.s3.eu-west-
1.amazonaws.com/tenzir-platform.yml \
  --parameters \
    ParameterKey=DomainName,ParameterValue=example.org \
    ParameterKey=UseExternalOIDC,ParameterValue=true \

ParameterKey=ExternalOIDCIssuerURL,ParameterValue=https://idp.exa
mple.org \
    ParameterKey=ExternalOIDCClientID,ParameterValue=<client-id>
\
    ParameterKey=ExternalOIDCClientSecret,ParameterValue=<client-
secret> \
    ParameterKey=ContainerImageUI,ParameterValue=<ui-image-uri> \
    ParameterKey=ContainerImageAPI,ParameterValue=<api-image-uri>
\
    ParameterKey=ContainerImageGateway,ParameterValue=<gateway-
image-uri> \
  --capabilities CAPABILITY_NAMED_IAM
```

— Runnable CloudFormation commands with exact parameters `product docs`

04 **CLI file loading and Pipelines-as-Code** NEW 95

Adds `-f` to the `tenzir` binary to load a pipeline definition from a file instead of inline. Supports Pipelines-as-Code by declaring pipelines under `tenzir.pipelines` in `<prefix>/etc/tenzir/tenzir.yaml`, with per-pipeline `name`, `description`, `definition`, `restart-on-error` (true/false/null or a duration like `1 minute`), `unstoppable` (run continuously, entering a failed state if it completes), `disabled` (keep the definition but don't start it), and `labels` (tags shown on the pipeline overview page at `app.tenzir.com`). Also supports running open, transformation-only pipelines from the CLI that auto-complete with stdin/stdout, with output format selectable via operators such as `write_csv`, `write_ssv`, and `write_parquet`.

Run a saved pipeline file instead of inlining the definition — useful for complex or version-controlled pipelines in CI.

```
$ tenzir -f pipeline.tql
```

Deploy a persistent Suricata ingest pipeline as code so it starts automatically with the node and restarts on failure.

```yaml
tenzir:
  pipelines:
    suricata-over-tcp:
      name: Onboard Suricata from TCP
      description: |
        Onboards Suricata EVE JSON from TCP port 34343.
      definition: |
        accept_tcp "0.0.0.0:34343" {
          read_suricata
        }
        publish "suricata"
      restart-on-error: 1 minute
      labels:
        - Suricata
        - Onboarding
      disabled: false
      unstoppable: true
```

Parse a YAML file from stdin, filter a field, and write the result as CSV to stdout — no source or sink operators needed.

```
$ cat input.yaml | tenzir 'read_yaml | select x | write_csv'
```

— All config fields and CLI flags named with runnable examples `product docs`

05 **Centralized TLS configuration in tenzir.yaml** NEW 94

Adds a `tenzir.tls` block in `tenzir.yaml` with `enable`, `skip-peer-verification`, `cacert`, `certfile`, `keyfile`, `password`, `tls-min-version`, and `tls-ciphers` to configure TLS once for all compatible operators (`accept_relp`, `accept_splunk`, `from_http`, `to_opensearch`, `accept_opensearch`, `to_splunk`). `tls-client-ca` and `require-client-cert` under `tenzir.tls` enable mutual TLS, rejecting clients without a certificate signed by the configured CA. A separate `plugins.platform` TLS block (`skip-peer-verification`, `cacert`, `certfile`, `keyfile`, `tls-min-version`, `tls-ciphers`) configures outbound TLS exclusively for the node-to-platform connection, overriding node-level settings. Documentation also names the node-level keys `tenzir.tls.min-version` (`"1.0"` – `"1.3"`) and `tenzir.tls.ciphers` (an OpenSSL cipher list) explicitly, covering `to_opensearch`, `to_splunk`, `accept_tcp`, `from_tcp`, `serve_tcp`, and `accept_opensearch`, and clarify that `tls-min-version` / `tls-ciphers` in `<configdir>/tenzir/plugin/platform.yaml` take precedence for Platform connections.

Enforce a minimum TLS 1.2 and a restricted cipher suite across all Tenzir operators in one place, rather than repeating TLS flags on every pipeline operator.

```yaml
tenzir:
  tls:
    tls-min-version: "1.2"
    tls-ciphers: "ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256"
    cacert: "/etc/ssl/certs/ca-certificates.crt"
```

Require clients to authenticate with a certificate when connecting to server-mode operators, enabling mutual TLS so unauthenticated clients are rejected.

```yaml
tenzir:
  tls:
    certfile: "/etc/tenzir/server.crt"
    keyfile: "/etc/tenzir/server.key"
    tls-client-ca: "/etc/tenzir/client-ca.crt"
    require-client-cert: true
```

Pin a specific client certificate and cipher list only for the node-to-platform connection, without affecting TLS settings for other operators.

```yaml
plugins:
  platform:
    cacert: "/etc/ssl/certs/ca-certificates.crt"
    certfile: "/etc/tenzir/platform-client.crt"
    keyfile: "/etc/tenzir/platform-client.key"
    tls-min-version: "1.2"
    tls-ciphers: "HIGH:!aNULL:!MD5"
```

Enforce TLS 1.2 minimum and restrict cipher suites for all node outbound and inbound connections (e.g. to Splunk, OpenSearch, or TCP listeners).

```yaml
tenzir:
  tls:
    min-version: "1.2"
    ciphers: "ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384"
```

Override TLS settings only for the Tenzir Platform connection, using a stricter cipher list than the node-wide default.

```yaml
tls-min-version: "1.3"
tls-ciphers: "ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256"
```

— Full set of TLS keys and precedence rules named with examples `product docs`

06 **Internal services config for Tenzir Platform** NEW 90

Adds environment variables for the Tenzir Platform's internal services: `TENZIR_PLATFORM_ADMIN_RULES` defines which users get admin access using rules like `auth_allow_all` or `auth_user` with a `user_id` (generate rules with `tenzir-platform tools print-auth-rule`); `TENZIR_PLATFORM_INTERNAL_AUTH_SECRET`, `TENZIR_PLATFORM_INTERNAL_TENANT_TOKEN_ENCRYPTION_KEY`, and `TENZIR_PLATFORM_INTERNAL_APP_API_KEY` set per-instance encryption and API secrets; `BLOB_STORAGE__NODES_PUBLIC_ENDPOINT_URL` lets nodes on separate networks reach S3-compatible blob storage via a different URL than the UI uses; `TENZIR_PLATFORM_USE_INTERNAL_WS_PROXY` routes UI frontend requests through the UI backend instead of directly to the Tenzir Gateway for zero-trust topologies; `PRIVATE_JWT_FROM_HEADER` accepts externally-supplied JWTs from an HTTP header (e.g. `X-Goog-IAP-JWT-Assertion`) for proxies like Google Cloud IAP; and `CORS_ALLOWED_ORIGINS` on the `websocket-gateway` service in `docker-compose.yaml` allows UI origins when the gateway and UI are on different domains.

Restrict platform admin rights to two specific Google OAuth2 users in a production deployment.

```text
TENZIR_PLATFORM_ADMIN_RULES=
[{"auth_fn":"auth_user","user_id":"google-oauth2:12345678901"},
{"auth_fn":"auth_user","user_id":"google-oauth2:12345678902"}]
```

Route node blob-storage traffic through an internal proxy when nodes cannot reach the public S3 endpoint directly.

```text
BLOB_STORAGE__NODES_PUBLIC_ENDPOINT_URL=https://s3.internal.example.com
```

Accept Google Cloud IAP JWTs from the standard assertion header so users behind IAP skip the OIDC login flow.

```env
PRIVATE_JWT_FROM_HEADER=X-Goog-IAP-JWT-Assertion
```

— Every env var named with example values, no limits given product docs

07 **Storage, batching, and catalog tuning config keys** NEW

90

Adds a set of `tenzir.yaml` keys to tune node performance: `tenzir.demand.max-batches` caps in-memory topic buffering for `publish` / `subscribe`; `tenzir.import.batch-size` (default 8,192) bounds events per table slice; `tenzir.import.batch-timeout` forwards buffered slices before they reach `batch-size`; `tenzir.import.read-timeout` (default 10 seconds) controls how long a read blocks before retrying; `tenzir.max-partition-size` (default 4Mi / 2^22 records) bounds partition size and catalog memory footprint; `tenzir.active-partition-timeout` (default 10 seconds) closes active partitions by time rather than record count; `tenzir.index.rules` defines per-field/per-type catalog sketch rules with `targets` and `fp-rate`, with `tenzir.index.default-fp-rate` setting the type-level default; and `tenzir.zstd-compression-level` tunes Zstd compression for partitions at rest.

Reduce catalog false-positive rate for a high-cardinality URL field and all IP fields to cut down irrelevant partition reads.

```yaml
tenzir:
  index:
    default-fp-rate: 0.01
    rules:
      - targets:
          - suricata.http.http.url
        fp-rate: 0.005
      - targets:
          - :ip
        fp-rate: 0.10
```

— Every config key and default named with one worked example product docs

08 **Node proxy, subprocess, and endpoint controls** NEW

88

Adds `tenzir.pipeline-subprocesses: true` (or `TENZIR_PIPELINE_SUBPROCESSES=true`) in `tenzir.yaml` to run pipeline execution in subprocesses for improved error resilience and resource utilization. Adds `tenzir.http-proxy`, `tenzir.https-proxy`, and `tenzir.no-proxy` to route

outbound operator traffic through an HTTP proxy with per-host/CIDR bypass. `tenzir.endpoint` can be set to `false` to disable node-to-node TCP connections (making the node reachable only via the Tenzir Platform) or to a custom address to override the default `127.0.0.1:5158` listener, including IPv6.

Route all outbound operator traffic through a corporate proxy while bypassing it for internal RFC-1918 ranges.

```yaml
tenzir:
  http-proxy: http://proxy.example.com:3128
  https-proxy: http://proxy.example.com:3128
  no-proxy: .internal,10.0.0.0/8
```

Enable pipeline subprocesses to isolate faults and improve resource utilisation without changing the binary invocation.

```yaml
tenzir:
  pipeline-subprocesses: true
```

Harden a cloud-only node by disabling direct TCP access so only the Tenzir Platform can reach it.

```yaml
tenzir:
  endpoint: false
```

Route all Tenzir outbound HTTPS and gRPC traffic through a corporate proxy, bypassing it for internal subnets.

```yaml
tenzir:
  http-proxy: http://proxy.example.com:3128
  https-proxy: http://user:password@proxy.example.com:3128
  no-proxy: .internal,10.0.0.0/8
```

— All config keys named with runnable YAML snippets product docs

09 **Pipeline parallelism controls** IMPROVED 85

Adds a `// parallelism: <n>` pipeline directive (accepting `disabled`, `max`, or a positive integer) to run CPU-bound operators across multiple cores, with a node-wide default settable via `tenzir.parallelism` in `tenzir.yaml`. `fuse=none`,

`fuse=all`, and `fuse=parallel` options control whether operators are fused into a single group, `limit_partitions=<n>` caps instances for partition-keyed operators like `summarize`, `group`, and `deduplicate`, and a `transform parallel { … }` block operator (with its own `fuse=<value>`) parallelizes just a CPU-bound subsection of a pipeline. Parallel execution was subsequently enabled by default for asynchronous transforms (increasing throughput for DNS, context, and Python workloads) and for competing-consumer sources (increasing throughput for AMQP and Google Cloud Pub/Sub workloads).

— Config directives named but no worked example given

Tenzir Node v6.18.0: Parallel processin…

---

10 **Config file discovery and per-plugin config files** NEW

68

Adds `--bare-mode` to disable all system- and user-specified configuration path lookups (for script-wrapping use cases) and `--config=/path/to/tenzir.yaml` to explicitly specify a config file path at startup. Also supports plugin-specific configuration files at `<configdir>/tenzir/plugin/<plugin>.yaml` as an alternative to inline `plugins.<name>` entries in `tenzir.yaml`.

— Flags and path named but no worked example product docs

THINNER COVERAGE BELOW

11 **Predefined dashboards in workspace.yaml** NEW 57

Adds a `dashboards` key under `workspaces.<id>` in `workspace.yaml` to pre-define dashboards for static workspaces, with named cells, a pipeline `definition`, a visualization `type` (table, bar, line, etc.), and grid coordinates (`x`, `y`, `w`, `h`).

— Config key named but no worked example provided product docs

---

12 **Legacy non-neo executor removed** BREAKING 43

The legacy non-neo executor has been removed; pipelines relying on the legacy executor must migrate to the neo executor.

— States the change but no migration steps given Tenzir Node v6.18.0: Parallel processin…

---

13 **Recursive search function** NEW 32

Adds a `search` function for recursive value matching across pipeline data.

— No parameters, syntax, or example given Tenzir Node v6.18.0: Parallel processin…

14 **Position-independent static Linux binaries** IMPROVED

25

Provides position-independent static Linux binaries.

— No build target or distribution details given Tenzir Node v6.18.0: Parallel processin…

15 **Experimental warning suppression via quiet** NEW 21

Adds experimental warning suppression via `quiet`.

— Bare mention, no syntax or scope given Tenzir Node v6.18.0: Parallel processin…

└──▷ **BREAKING ON UPGRADE**

! The legacy non-neo executor has been removed; pipelines relying on the legacy executor must migrate to the neo executor.

Tenzir Node v6.18.0: Parallel processing across transforms and sources (2026-09-22)

BLOG ↗ LATEST STABLE

Tenzir v6.18.0 adds parallel pipeline execution, a recursive `search` function, and restores SentinelOne Data Lake operators.

Adds a `search` function for recursive value matching across pipeline data.

Adds experimental warning suppression via `quiet`.

Enables parallel execution for asynchronous transforms, increasing throughput for DNS, context, and Python workloads.

Enables parallel competing-consumer sources, increasing throughput for AMQP and Google Cloud Pub/Sub workloads.

Provides position-independent static Linux binaries.

└──▷ **BREAKING ON UPGRADE**

! The legacy non-neo executor has been removed; pipelines relying on the legacy executor must migrate to the neo executor.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

runZero is an asset discovery and inventory platform that identifies and maps all devices on networks without requiring agents.

> runZero's biggest window additions are a vulnerability suppression/unsuppression API, an OCSF 1.9.0 export and a new Cryptographic Bill of Materials (CBOM) export, alongside OAuth support for MCP, more granular custom roles and API-client permissions, and a round of scripting, reporting, and performance improvements.

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

01 **OCSF 1.9.0 export from asset inventory** NEW 80

Adds an OCSF 1.9.0 export from the asset inventory, exporting assets as Device Inventory Info events and observed services as Live Evidence Info events, available from the inventory export menu and the Export API. Fields without an OCSF equivalent are preserved under the event's `unmapped` object.

— Names event types, access points, and the unmapped field explicitly.

changelog-20260922-3fd0bc65

02 **Custom integration scripting enhancements** NEW 75

Enables running custom integration scripts on console instances in addition to Explorers and the CLI. Adds an in-product script authoring and testing environment using the AI model of your choice. Scripts now support settings shared across protocols like HTTP and TLS. Allows users to submit scripts for inclusion as native runZero sources.

— Groups four related scripting additions naming each surface.

changelog-20260922-3fd0bc65

03 **Custom roles and granular permissions across API keys, users, groups, and SSO**

NEW 70

Supports assigning custom roles to API keys, default user roles, group permissions, SSO defaults, and MCP connections. API clients can now be created with custom permissions capped by the creating user's own access rights, and account-level API clients can use granular authentication scopes in addition to roles. The list of default roles selectable in SSO settings now includes the Annotator role.

— Names every surface gaining custom roles but no config steps.

changelog-20260922-3fd0bc65

04 **Cryptographic Bill of Materials (CBOM) export** NEW

65

Adds a Cryptographic Bill of Materials (CBOM) as a new CycloneDX JSON export and detailed report, letting users identify weak or deprecated algorithms across the asset inventory.

Generate a Cryptographic Bill of Materials for an org to identify weak or deprecated algorithms across the asset inventory.

> 📍 In the runZero console, go to Reports, select `Cryptographic Bill of Materials (CBOM)`, then choose `Export as CycloneDX JSON` to download the artifact.

— Comes with exact UI steps to generate the export. changelog-20260922-3fd0bc65

05 **OAuth support for MCP service** NEW

65

Adds OAuth support to the MCP service, including the ability to use Export Tokens and Organization Keys as OAuth Client Secrets, in addition to API Clients.

— Names the credential types but no setup steps. changelog-20260922-3fd0bc65

THINNER COVERAGE BELOW

06 **Vulnerability suppression and unsuppression via API** NEW

55

Adds vulnerability suppression and unsuppression via the runZero API, covering vulnerabilities, vulnerability groups, and findings.

— Names affected object types but not the exact endpoint. changelog-20260922-3fd0bc65

07 **AI-generated content labeling** NEW

55

runZero now labels AI Threads and AI Artifacts as AI-generated, embedding the label in exported reports, downloaded files, and generated custom-integration scripts.

— Explains where the label propagates but is automatic. changelog-20260922-3fd0bc65

08 **Promiscuous ARP probe toggle on scan page** NEW

50

Adds a configuration option on the scan page to toggle the promiscuous ARP probe.

— Thin description but gives an exact UI location. changelog-20260922-3fd0bc65

09 **Reports page and External Reporting enhancements** IMPROVED

50

The reports page now displays a paginated, searchable list of all analysis task-based reports with delete support. Users can now share standard reports through the External Reporting feature.

— Clear UI location but no deeper mechanism given. changelog-20260922-3fd0bc65

10 **Dashboards column in query library** NEW 50

Adds a Dashboards column to the query library showing how many personal dashboards use a saved query and whether it appears on other dashboards.

— Describes exact new column and what it shows. changelog-20260922-3fd0bc65

11 **WS-Discovery XAddr contributes to asset hostnames** IMPROVED 50

WS-Discovery XAddr endpoints now contribute to asset hostnames.

— Specific mechanism named but automatic, no user action. changelog-20260922-3fd0bc65

12 **Issue workflow improvements** IMPROVED 50

Adds the ability to close an Issue directly from the Issue Details page. Issue owners now receive an email with a direct link when assigned an Issue.

— Two thin but distinct Issues workflow changes named together.

changelog-20260922-3fd0bc65

13 **Task management UI improvements** IMPROVED 50

Users can now reprocess tasks directly from the user interface. Task summaries now indicate when a task involved more than 1,000 asset changes.

— Names the threshold and the new UI action. changelog-20260922-3fd0bc65

14 **Performance improvements across scanning, tasks, and exports** IMPROVED 45

runZero optimized MAC merging so large scan jobs now process faster, now processes directory and device data concurrently to improve task performance, sped up vulnerability exports, and no longer counts background MCP connections against API usage.

— Names four improved areas but no metrics or user action. changelog-20260922-3fd0bc65

15 **Mermaid diagrams in AI Threads** NEW 40

Adds Mermaid-based charting and diagram support to the AI Threads feature.

— Names the feature and library but no usage detail. changelog-20260922-3fd0bc65

16 **Re-invite existing external users** IMPROVED 40

Users can now re-invite an external user who already has access to the account.

— Bare one-line UI addition. changelog-20260922-3fd0bc65

17 **Sort Findings by name** NEW 35

Adds the ability to sort Findings by name.

— Bare one-line UI addition. changelog-20260922-3fd0bc65

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

runZero adds vulnerability suppression API, OCSF 1.9.0 export, CBOM CycloneDX JSON, OAuth for MCP, and custom API key roles.

└──▷ **HOW TO FIND IT**

Generate a Cryptographic Bill of Materials for an org to identify weak or deprecated algorithms across the asset inventory.

> 📍 In the runZero console, go to Reports, select Cryptographic Bill of Materials (CBOM), then choose Export as CycloneDX JSON to download the artifact.

Adds vulnerability suppression and unsuppression via the runZero API, covering vulnerabilities, vulnerability groups, and findings.

Adds an OCSF 1.9.0 export from the asset inventory, exporting assets as Device Inventory Info events and observed services as Live Evidence Info events, available from the inventory export menu and the Export API; fields without an OCSF equivalent are preserved under the event's `unmapped` object.

Adds a Cryptographic Bill of Materials (CBOM) as a new CycloneDX JSON export and detailed report.

Adds OAuth support to the MCP service, including the ability to use Export Tokens and Organization Keys as OAuth Client Secrets, in addition to API Clients.

Supports assigning custom roles to API keys, default user roles, group permissions, SSO defaults, and MCP connections.

+23 more

WAS THIS USEFUL? 👍 👎

Censys

Feedback ⓘ September 21, 2026 (2026-09-21) BLOG ↗ RANK

Censys is a search engine for internet-connected devices that indexes servers, certificates, and host data for security research and threat intelligence.

> Censys shipped a new ServiceNow TISC integration for enriching, rescanning, and pivoting on assets directly from ServiceNow, alongside 16 new platform fingerprints and two new ASM risk fingerprints for exposed management interfaces and a Cisco vulnerability.

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

01 New ASM risk fingerprints for exposed and vulnerable assets

NEW 88

Adds two new ASM risk fingerprints, queryable via the `risks.name` field: `"Exposed Check Point Security Management and Log Server"` to find internet-exposed management-plane assets, and `"Vulnerable Cisco Secure Email Gateway [CVE-2026-76461]"` to identify assets affected by the actively-exploited SQL injection vulnerability CVE-2026-76461.

Find all internet-exposed Check Point Security Management or Log Servers in your ASM workspace to prioritize remediation of management-plane exposure.

```
$ risks.name: "Exposed Check Point Security Management and Log Server"
```

Identify assets running a version of Cisco Secure Email Gateway affected by the actively-exploited SQL injection vulnerability CVE-2026-76461.

```
$ risks.name: "Vulnerable Cisco Secure Email Gateway [CVE-2026-76461]"
```

— Exact risk names and runnable risks.name query provided September 21, 2026 (2026-09-21)

02 ServiceNow TISC integration NEW

64

Adds a ServiceNow TISC integration that can automatically or manually enrich IPs, web properties, and certificates in ServiceNow TISC with Censys Platform data. From within ServiceNow TISC, users can also initiate a Censys rescan of a host or web property, retrieve scan history for an IP address, and initiate a CensEye automated pivoting job to find assets related to an IP, web property, or certificate.

— Names integration and four capabilities but no config keys or endpoints

September 21, 2026 (2026-09-21)

THINNER COVERAGE BELOW

03 **16 new platform fingerprints** NEW 53

Adds 16 new platform fingerprints covering Check Point Security Management and Log Server, Apache CloudStack, Apache Cocoon, Cisco Secure Email Gateway, Citrix XenMobile Server, Docker Registry Browser, GeoVision GeoWebServer, IBM WebSphere Liberty, Lucee, OpenMetadata, phpList, Plone, PrivateBin, pypiserver, Roadiz CMS, and TYPO3.

— Names all 16 fingerprints but no query mechanism given September 21, 2026 (2026-09-21)

September 21, 2026 (2026-09-21) BLOG ↗ LATEST STABLE

Censys adds ServiceNow TISC integration, 16 new platform fingerprints, and two new ASM risk fingerprints.

└──▷ **TRY IT**

Find all internet-exposed Check Point Security Management or Log Servers in your ASM workspace to prioritize remediation of management-plane exposure.

```
$ risks.name: "Exposed Check Point Security Management and Log Server"
```

Identify assets running a version of Cisco Secure Email Gateway affected by the actively-exploited SQL injection vulnerability CVE-2026-76461.

```
$ risks.name: "Vulnerable Cisco Secure Email Gateway [CVE-2026-76461]"
```

Adds a ServiceNow TISC integration to automatically or manually enrich IPs, web properties, and certificates in ServiceNow TISC with Censys Platform data.

Adds the ability to initiate a Censys rescan of a host or web property directly from ServiceNow TISC.

Adds the ability to retrieve scan history for an IP address from within ServiceNow TISC.

Adds the ability to initiate a CensEye automated pivoting job from ServiceNow TISC to find assets related to an IP, web property, or certificate.

Adds 16 new platform fingerprints covering Check Point Security Management and Log Server, Apache CloudStack, Apache Cocoon, Cisco Secure Email Gateway, Citrix XenMobile Server, Docker Registry Browser, GeoVision GeoWebServer, IBM

> WebSphere Liberty, Lucee, OpenMetadata, phpList, Plone, PrivateBin, pypiserver, Roadiz CMS, and TYPO3.

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

defectdojo

Feedback ⓘ 3.3.200 NOTES ↗ RANK

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo's main addition this window is an asset hierarchy feature that suggests `deploys to` relationships from container image evidence, backed by scanners now emitting image locations into the Locations model; a smaller performance fix debounces product grade recalculation.

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

01 Suggested deploy-to edges from container evidence NEW

79

When Container Image Locations are enabled, DefectDojo detects undrawn `deploys to` relationships by matching image repositories across assets and surfaces each pair as a reviewable suggestion on the asset hierarchy page, backed by container-image evidence. Accepting a suggestion draws the `deploys to` edge from the image-building asset to the running asset (the deployed asset inherits its host's exposure) and the accepted edge carries an origin of container evidence to distinguish it from manually drawn or connector-declared edges; dismissing a suggestion suppresses that asset pair (further images for the pair are counted but do not reopen it, and dismissed suggestions can be reopened from the same dialog). Reviewing suggestions requires the Asset Hierarchy view permission, while accepting or dismissing requires the edit permission plus edit access to both assets.

Review and accept container-derived deploy relationships on the asset hierarchy page to propagate exposure from host assets to deployed assets automatically.

> 📍 In the Asset Hierarchy page, look for the banner counting open suggestions, click Review to open the list, then click Accept on a pair to draw the deploys-to edge, or Dismiss to suppress it.

— Full mechanism, permissions, and UI path given; no exact endpoint/flag. product docs

THINNER COVERAGE BELOW

02 Container image location emission from scanners NEW

46

Container scanners now emit the scanned image as an Image location, linking container images to their repository in the Locations model. This is the underlying data source that powers the asset hierarchy's deploy-to suggestions.

— Names the `Image` location and Locations model but no usage steps. 3.3.200

03 **Debounced product grade recalculation** IMPROVED 24

Product grade recalculation is now debounced per product to improve performance on large imports.

— One-line performance note with no mechanism or numbers. 3.3.200

3.3.200 NOTES ↗ LATEST STABLE

DefectDojo 3.3.200 adds container image location emission from container scanners and debounced product grade recalculation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.3.200 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.3.200
```

Container scanners now emit the scanned image as an `Image` location, linking container images to their repository in the Locations model.

Debounces product grade recalculation per product to improve performance on large imports.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Okta for AI Agents

Feedback ⓘ changelog · Sep 22, 2026

CHANGELOG ↗ RANK

> Okta introduced a full AI Agent identity platform — Agent Gateway runtime enforcement, Agent SSO, shadow AI discovery, ephemeral credentialing, HITL guardrails, ISPM, and FedRAMP/DoD/HIPAA compliance tiers — alongside the 2026.09.0 platform release adding MCP server DCR/OIN registration, new malware-proxy and remote-desktop network zone categories, NFC and Linux authenticators, and a wide set of admin console, provisioning, and limit improvements.

└──▷ **WHAT SHIPPED · 29 FEATURES** most completely described first ? what's the number?

01 New IP service categories in Enhanced Dynamic Network Zones

NEW 95

Adds `REMOTE_DESKTOP` IP service category to Enhanced Dynamic Network Zones, letting admins include or exclude remote-desktop traffic in global session and app sign-in policies for precise access control. Adds `MALWARE_PROXIES` IP service category, powered by Okta CyberDefense, covering proxy services tied to malware and botnet activity including 911 S5, NSOCKS, iProxy, and BHProxies.

Block access from known malware proxy networks by adding the `MALWARE_PROXIES` category to an Enhanced Dynamic Network Zone and denying it in the global session policy.

> 📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add `MALWARE_PROXIES` to the excluded IP service categories, then go to Security › Global Session Policy and add a rule that denies access when the network zone matches.

Block remote-desktop originating traffic at the app sign-in policy level using the new `REMOTE_DESKTOP` IP service category.

> 📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add `REMOTE_DESKTOP` to the excluded IP service categories, then open the target app's Sign-on Policy and add a rule denying access when the zone matches.

— Exact category names and admin console steps provided `changelog-20260922-2fd4fb13`

02 MCP server registration via OIN catalog and DCR NEW

80

Supports Dynamic Client Registration (DCR) when manually registering an MCP server so Okta automatically registers a client with the provider and populates credentials. MCP servers can also be added directly from an Okta Integration Network (OIN) catalog entry, eliminating manual connection-detail entry, and tool discovery at registration time lets admins test credentials and view available tools before completing setup. The MCP server registration UI was updated with improved scope handling and tool visibility.

Register an MCP server from the OIN catalog to avoid manually entering connection details and immediately discover its available tools.

> 📍 In the Admin Console, go to Applications › MCP Servers, click `Add MCP Server`, select `Add from OIN catalog`, locate the provider entry, choose Dynamic Client Registration (DCR) as the registration method, and click `Test credentials and discover tools` to verify the connection.

— Concrete UI flow given, all constituent additions named `changelog-20260922-2fd4fb13`

---

03 **Failed WebAuthn enrollment logging in System Log** NEW

80

Logs failed WebAuthn (FIDO2) enrollment attempts in the System Log under the `user.mfa.factor.activate` event, with debug data including `AAGUID`, `isBackupEligible`, and matched authenticator groups.

— Names exact event and fields to search System Log `changelog-20260922-2fd4fb13`

---

04 **Okta Integration Wizard for custom app integrations** NEW

70

Introduces the Okta Integration Wizard (OIW) to create and deploy custom app integrations configuring SSO, SCIM provisioning, Entitlement Management, Universal Logout, API service integration, and custom API Integration Actions — reusable as a template across multiple app instances.

— Names configurable surfaces, no exact navigation given `changelog-20260922-2fd4fb13`

---

05 **FedRAMP, DoD IL5, and HIPAA compliance tiers** NEW 65

The Okta for AI Agents Core SKU supports FedRAMP Moderate/High and DoD Impact Level 5 environments, covering agent registration, cryptographic authentication, lifecycle provisioning, and SIEM-integrated audit logging. The full Okta for AI Agents SKU adds HIPAA support, including ISPM and Okta Privileged Access (OPA) capabilities.

— Names SKUs and compliance scopes, no purchase/config steps `product docs`

---

06 **NFC authenticator for frontline workers** NEW 65

Adds NFC authenticator support for frontline workers on Windows shared workstations — users tap an NFC badge and enter a PIN to meet MFA requirements without a phone,

password, or shared account.

— Describes exact user flow, no admin setup steps given changelog-20260922-2fd4fb13

07 **Unified identity management for AI agents** NEW 65

Admins can set an external identifier to enforce uniqueness, configure creation and matching criteria for imported AI agents, and preview imports before registration. The AI agent Client Registration tab now displays authentication methods vertically with a Configure button linking to per-method configuration pages.

— Names UI element and button, no full workflow steps changelog-20260922-2fd4fb13

08 **Low-Code Sign In Customization** NEW 65

Lets admins modify text, colors, and images in the Sign-In Widget and Interstitial authentication using a JSON templating language with syntax highlighting and suggestions — no code editor required.

— Names templating language and UI targets, no exact menu path

changelog-20260922-2fd4fb13

09 **Entitlement import safeguards for app provisioning** NEW 65

Adds entitlement import safeguards configurable per app using percentage-based or absolute count thresholds to prevent accidental removal of app roles or licenses on user unassignment, with an option to block imports that modify or delete entitlement schemas.

— Names threshold types, no exact config field path changelog-20260922-2fd4fb13

10 **Higher Access Request configuration limits** IMPROVED 65

Increases Access Request limits: users per task or question to 25 (from 10), entitlement bundles in an access level condition to 1,000 (from 100), groups in an access level condition to 1,000 (from 500), request type configuration lists per org to 250 (from 100), and request types per org to 750 (from 500).

— All limit numbers named, no config path given changelog-20260922-2fd4fb13

11 **Shadow AI discovery and Universal Directory registration** NEW

60

Multi-surface shadow AI discovery continuously scans browsers, managed endpoints, and cloud environments to uncover unvetted agents and unauthorized OAuth consent grants; discovered agents are then registered into Universal Directory alongside workforce users, with explicit human ownership, business purpose, and cryptographic credentials assigned.

— Explains scanning surfaces and registration fields, no UI steps `product docs`

12 **Prompt-level MCP tool scoping** NEW

60

Filters unauthorized MCP tools based on the delegating user's identity before the model executes, reducing schema token overhead by over 90%.

— Has concrete metric but no setup steps `product docs`

13 **Okta Verify for Linux desktop** NEW

60

Extends Okta Verify with official Linux desktop support (Ubuntu 24.04 and 26.04), enabling phishing-resistant, hardware-backed Okta FastPass for Linux users.

— Named OS versions, no install/config steps `changelog-20260922-2fd4fb13`

14 **HA and incremental imports for Generic Database connector** IMPROVED

60

Supports high availability for the On-prem Connector for Generic Databases by allowing multiple Okta On-Premises SCIM Agents to be assigned to a single app instance, and adds incremental imports that retrieve only users and entitlement assignments changed since the last successful import.

— Names connector and mechanism, no config steps `changelog-20260922-2fd4fb13`

15 **Applications page filtering, search, and export** IMPROVED

60

Enhances the Applications page with filtering by type and status, search by name or client ID, a last-modified-date view, and CSV export for audit-ready reports.

— Names all UI additions, no exact navigation path `changelog-20260922-2fd4fb13`

16 **Okta On-prem SCIM agent replaces SCIM Server agent** BREAKING

60

Replaces Okta On-prem SCIM Server agent with Okta On-prem SCIM agent, reducing dependencies and enabling new features; existing deployments referencing the old agent name must migrate to the new Okta On-prem SCIM agent.

— Names old and new agent, notes migration need, no exact migration steps

changelog-20260922-2fd4fb13

THINNER COVERAGE BELOW

17 **Agent Gateway for AI agent identity verification** NEW

55

Agent Gateway verifies agent identity, brokers credentials at runtime, and logs every tool call, all without requiring code changes to agents or tools.

— Describes mechanism but no config surface or UI path product docs

18 **Ephemeral secretless credentialing for agents** NEW

50

Issues ephemeral, secretless tokens and dynamically vaults credentials to eliminate static API keys and standing administrative access for AI agents.

— Clear mechanism, no config keys or commands given product docs

19 **Human-in-the-loop guardrails for high-risk actions** NEW

50

Requires step-up human approval before high-risk agent actions such as financial transactions, bulk database exports, or permission changes.

— Names risk categories, no configuration path product docs

20 **AD/LDAP agent disruption and recovery email notifications**

NEW

50

Adds email notifications for Active Directory and LDAP agent disruption and recovery, configurable in the Admin Console.

— Names configuration location, lacks exact setting path changelog-20260922-2fd4fb13

21 **New provisioning integrations: WordPress, Ivanti, Chef, Vercel, Sophos**

NEW

50

Adds provisioning support for WordPress, Ivanti, Progress Chef, Vercel, and Sophos Cloud integrations.

— Names all five integrations, no setup detail changelog-20260922-2fd4fb13

22 **Agent SSO for Cross-App Access agents** NEW

45

Agent Single Sign-On manages Cross-App Access agents within core SSO, with centralized policies and full lifecycle governance.

— Named capability but no configuration steps given product docs

23 **Programmatic kill switch for agent access** NEW 45

Enforces a programmatic kill switch to immediately revoke an agent's access, producing a full audit trail when unexpected behavior is detected.

– Describes behavior, no API or command named `product docs`

24 **Improved DBSSO device probing** IMPROVED 45

Improves DBSSO device probing to use device-registered conditions from an app's sign-on policy instead of org-wide probing methods.

– Names mechanism change, no configuration steps `changelog-20260922-2fd4fb13`

25 **Identity Security Posture Management for agents** NEW 40

Identity Security Posture Management (ISPM) evaluates the security posture of each agent identity across users, apps, and permissions.

– Names capability, no metrics or workflow detail `product docs`

26 **Automated agent access certification campaigns** NEW 40

Schedules automated access certification campaigns to continuously validate and revoke stale agent permissions as projects evolve.

– Brief description, no scheduling config shown `product docs`

27 **Access governance for MLOps supply chain** NEW 35

Governs access to model registries, code repositories, and development environments to secure the broader MLOps supply chain.

– Names surfaces covered but no mechanism detail `product docs`

28 **Desktop MFA Factor Discovery for Windows** IMPROVED 35

Shows Windows users only their enrolled MFA factors instead of a fixed list of all available factors.

– Brief behavioral description, no config path `changelog-20260922-2fd4fb13`

29 **Higher realm assignment limit per profile source** IMPROVED 35

Increases the maximum realm assignments per profile source from 30 to 100.

– Single number change, minimal context `changelog-20260922-2fd4fb13`

└──▷ **BREAKING ON UPGRADE**

! Okta On-prem SCIM Server agent has been replaced by Okta On-prem SCIM agent — existing deployments referencing the old agent name must migrate to the new Okta On-

prem SCIM agent.

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

Okta 2026.09.0 adds AI agent identity management, MCP server DCR and OIN catalog support, NFC authenticator, Okta Verify for Linux, and malware proxy detection.

└──▷ **HOW TO FIND IT**

Block access from known malware proxy networks by adding the `MALWARE_PROXIES` category to an Enhanced Dynamic Network Zone and denying it in the global session policy.

> 📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add `MALWARE_PROXIES` to the excluded IP service categories, then go to Security › Global Session Policy and add a rule that denies access when the network zone matches.

Block remote-desktop originating traffic at the app sign-in policy level using the new `REMOTE_DESKTOP` IP service category.

> 📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add `REMOTE_DESKTOP` to the excluded IP service categories, then open the target app's Sign-on Policy and add a rule denying access when the zone matches.

Register an MCP server from the OIN catalog to avoid manually entering connection details and immediately discover its available tools.

> 📍 In the Admin Console, go to Applications › MCP Servers, click `Add MCP Server`, select `Add from OIN catalog`, locate the provider entry, choose Dynamic Client Registration (DCR) as the registration method, and click `Test credentials and discover tools` to verify the connection.

Adds `REMOTE_DESKTOP` IP service category to Enhanced Dynamic Network Zones, letting admins include or exclude remote-desktop traffic in global session and app sign-in policies for precise access control.

Adds `MALWARE_PROXIES` IP service category to Enhanced Dynamic Network Zones, powered by Okta CyberDefense, covering proxy services tied to malware and botnet activity (including 911 S5, NSOCKS, iProxy, and BHProxies).

Supports Dynamic Client Registration (DCR) when manually registering an MCP server — Okta automatically registers a client with the provider and populates credentials.

Enables adding MCP servers directly from an Okta Integration Network catalog entry, eliminating the need to enter connection details manually.

Adds tool discovery for MCP servers at registration time, letting admins test credentials and view available tools before completing setup.

+19 more

└─▷ **BREAKING ON UPGRADE**

! Okta On-prem SCIM Server agent has been replaced by Okta On-prem SCIM agent — existing deployments referencing the old agent name must migrate to the new Okta On-prem SCIM agent.

WAS THIS USEFUL?  

iDenfy

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

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

> iDenfy's Sprint 135 release adds Bank Card Verification and a Blocklist API endpoint, extends KYB with field filtering and shared identity verification across companies, adds native-script CSV exports, automatic branding import from your website, and document-capture improvements including real-time framing feedback and expanded injection detection, while raising the iOS SDK's minimum deployment target to iOS 15.

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

01 **Field filtering for KYB company endpoint** NEW 70

Adds a `fields` parameter to the KYB company endpoint to return only the data fields you need.

– Endpoint and parameter named, but no example values given. changelog-20260922-b41dfb9e

02 **Native-script data in KYC CSV export** NEW 60

Adds an Original language data section to KYC CSV reports, exporting document values in their native script as separate columns.

– Names exact report section and export format. changelog-20260922-b41dfb9e

THINNER COVERAGE BELOW

03 **iOS SDK minimum deployment target raised** BREAKING 55

The iOS SDK minimum deployment target is raised to iOS 15, a breaking change for apps still targeting earlier iOS versions.

– Clear version threshold gives a migration action. changelog-20260922-b41dfb9e

04 **Automatic branding import from website** IMPROVED 45

Supports PNG logo uploads and enables branding themes to be prefilled from your website, automatically pulling in colors and logo.

– Describes source and effect, no config path given. changelog-20260922-b41dfb9e

05 **Document capture feedback and injection detection expansion** IMPROVED 45

Adds real-time framing feedback during document capture, and extends injection detection to images that did not go through the iDenfy capture flow (experimental).

– Two capture-flow additions named, mechanisms not detailed. changelog-20260922-b41dfb9e

06 **iOS SDK badge and iOS 27 support** IMPROVED 45

Adds a Recommended document badge and iOS 27 support in the iOS SDK (v9.2.2).

– SDK version and OS support named, behaviour unclear. changelog-20260922-b41dfb9e

07 **Blocklist API endpoint** NEW 35

Adds a Blocklist API endpoint.

– No path, method, or fields specified. changelog-20260922-b41dfb9e

08 **KYB shared identity verification across companies** NEW 30

Adds KYB shared identity verification across companies.

– Sharing mechanism between companies not detailed. changelog-20260922-b41dfb9e

09 **Bank Card Verification** NEW 25

Adds Bank Card Verification capability.

– No mechanism, scope or workflow described. changelog-20260922-b41dfb9e

└─▷ BREAKING ON UPGRADE

! The iOS SDK minimum deployment target is raised to iOS 15.

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

iDenfy adds Bank Card Verification, Blocklist API, KYB fields filtering, CSV native-script export, and branding auto-import in Sprint 135

Adds a `fields` parameter to the KYB company endpoint to return only the data fields you need.

Adds a Blocklist API endpoint.

Adds Bank Card Verification capability.

Adds KYB shared identity verification across companies.

Adds an Original language data section to KYC CSV reports, exporting document values in their native script as separate columns.

+6 more

└─▷ BREAKING ON UPGRADE

! The iOS SDK minimum deployment target is raised to iOS 15.

WAS THIS USEFUL? 👍 👎

Didit

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

Didit Web SDK 0.4.0 brings native camera-control options to the browser, including rear-camera and switcher-hide flags.

└──▷ **TRY IT**

Open a verification session on the rear camera with the switcher hidden when linking users directly to the verification URL instead of using the SDK.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

Adds `defaultDocumentCamera` and `defaultLivenessCamera` SDK options to the Web SDK, allowing the document and face-capture steps to open on the rear camera by default (e.g. for kiosk or operator-assisted flows).

Adds `showDocumentCameraSwitchButton` and `showLivenessCameraSwitchButton` SDK options to the Web SDK to hide the in-capture camera switcher and lock users to the chosen lens.

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗

RANK

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

> ConductorOne's biggest changes this window are infrastructure-level: EU data residency, a public REST API for AI governance objects, Terraform support for MCP sources, and a breaking split of credential revocation from secret deletion — alongside persistent AI-generated reports, AI-assisted task diagnosis, and a string of access-review and campaign refinements.

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

01 Credential revoke decoupled from secret deletion

BREAKING

90

`POST /api/v1/app-secrets-admin/revoke` now queues a credential revoke separately from secret deletion, while `DELETE /api/v1/secrets/{vault_id}` only deletes the secret and no longer revokes its credential. Service SDKs rename their Revoke method to Delete to match the new behavior.

Revoke a credential after deleting a shared secret — now two separate calls since deletion no longer implies revocation.

```
$ curl -X DELETE
'https://<tenant>.conductor.one/api/v1/secrets/<vault_id>' \
  -H 'Authorization: Bearer <token>'

curl -X POST 'https://<tenant>.conductor.one/api/v1/app-secrets-
admin/revoke' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vault_id": "<vault_id>"}'
```

— Exact endpoints, methods and a working curl example given. changelog-20260922-e7a9b8d7

02 Spend-remedy task type and endpoint removal

BREAKING

90

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

— Exact endpoint, constant value and required client changes named.

changelog-20260922-e7a9b8d7

03 **Saveable and reusable C1AI reports** NEW 80

A Save report action in the AI assistant names and persists a report so it isn't lost between sessions. A new `Platform > Reporting` page lists all saved reports with rename, delete, refresh, and Edit with Report assistant controls, and each saved report has a Function tab to inspect the underlying logic and data sources behind its results.

Find and re-run a previously saved access report without hunting through conversation history.

📍 Navigate to Platform › Reporting, locate the saved report by name, then select Refresh to re-run it against the latest data, or Edit with Report assistant to adjust its parameters in a conversation.

— Exact UI navigation path and named controls given. changelog-20260922-e7a9b8d7

04 **AI-assisted stuck task diagnosis** NEW 75

Super Admins can ask C1AI to diagnose a stuck, failed, or waiting task (e.g. 'why is task #4821 stuck?'); the assistant traces the task's approval and provisioning state, access-profile origin, app-account access, and connector handoff, then recommends whether to approve, deny, or retry, without taking disruptive action itself.

Diagnose a stuck provisioning task as a Super Admin without touching any live state.

📍 In the C1 console, open the stuck task's detail page, then select Ask C1AI and type: 'Why is this task stuck and what should I do next?'

Diagnose a task that is stuck in provisioning without taking any disruptive action — useful before escalating or retrying.

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

— Exact prompt example and traced diagnostic scope given. changelog-20260922-e7a9b8d7

05 **EU data residency instance** NEW 65

A new EU data residency instance is selectable at domain registration, affecting all tenant-specific URLs including login, redirect URIs, API and webhook endpoints, MCP server, and CLI configuration.

— Lists affected surfaces but not the exact selection flow. changelog-20260922-e7a9b8d7

06 **Change history and justification for configuration changes**

NEW 65

A Change history icon on configuration object detail pages (connectors, policies, entitlement routing rules, SSO, branding settings, AI Gateway spend controls) shows a before-and-after timeline of every change. Connector configuration changes can now optionally require a justification, captured alongside that history.

— Named object types but only a UI icon as entry point. changelog-20260922-e7a9b8d7

07 **Set assignee step in Transformation rules** NEW 65

A new Set assignee step in Automation > Transformation rules automatically assigns a matched access-review finding to a specific person, overriding any automatic assignee for those findings.

Automatically assign access-review findings to the right owner using a Transformation rule instead of manual triage.

> 📍 In the C1 console, go to Automation › Transformation rules, open or create a rule, add a new step, select Set assignee, and specify the target user — this assignment will override any automatic assignee for matched findings.

— Exact menu path and override behavior are given. changelog-20260922-e7a9b8d7

08 **Access review and campaign management improvements** IMPROVED 65

Access-review revoke-task reviewer denial reasons are now automatically included in ticketing system templates, with new template variables available for custom templates. Campaign scope now expands to include users with pending (unaccepted) account invitations. Conflict monitors can handle groups of very different sizes (e.g., 90 and 11) as long as the combined total stays under 1,024 members. Access-review tables also gain a Resource description column, showing custom descriptions or connector-synced metadata as a fallback.

— Four concrete constituents named, but no single actionable path.

changelog-20260922-e7a9b8d7

09 **Public REST API for AI governance objects** NEW 60

Classifiers, tool gates, classifier bindings, templates, and the agent classifier policy are now exposed through the public REST API, enabling programmatic management of AI governance configuration.

— Named objects listed but no specific endpoint paths given. changelog-20260922-e7a9b8d7

THINNER COVERAGE BELOW

10 **Terraform support for MCP sources** NEW 55

New Terraform resources let admins manage hosted and external MCP sources as code via C1's Terraform provider, covering source details, authentication, and tool enablement policy.

– Names the surface but not specific resource names or syntax.

changelog-20260922-e7a9b8d7

11 **Apps and Entitlements tabs in My resources** NEW 50

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

– UI path named but no further mechanism described. changelog-20260922-e7a9b8d7

12 **Explicit toolset grant required for MCP access** BREAKING 50

MCP access control now requires an explicit toolset grant on the MCP server — app membership or an admin role alone no longer implies access, so existing setups relying on implicit access must be updated.

– Explains the behavior change but no migration steps given. changelog-20260922-e7a9b8d7

13 **Safer handling of legacy name-only revoke criteria** BREAKING 50

A revoke criteria scoped only by a legacy resource-type name now fails the step instead of silently matching every app sharing that name; automations still relying on name-only matching must be re-scoped.

– Behavior change explained with a remediation hint, no example.

changelog-20260922-e7a9b8d7

14 **Unified MCP integration setup flow** IMPROVED 45

Registered MCP integrations, both catalog and external, are now unified into a single guided setup flow covering connection, access, tool approval, and authentication.

– Describes flow stages but not exact navigation. changelog-20260922-e7a9b8d7

15 **Inline entitlement request forms in Slack** NEW 45

Entitlement request forms now render inline in the Slack modal when requesting or extending access via the C1 app for Slack.

– Names the surface but no further mechanism detail. changelog-20260922-e7a9b8d7

16 **Agent-to-agent discovery over MCP** NEW 40

Customer agents connected to a C1 MCP endpoint can now discover and converse with c1aw agents, enabling agent-to-agent interaction through the MCP integration.

— Brief description with no setup steps or scope detail. `changelog-20260922-e7a9b8d7`

└──▷ **ALSO FROM THESE RELEASES**

.
.
.
.
.
.

└──▷ **BREAKING ON UPGRADE**

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

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

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

changelog · Sep 22, 2026 | CHANGELOG ↗ | LATEST | STABLE

ConductorOne adds EU data residency, saveable AI reports, REST API for AI governance, MCP Terraform support, and a `POST /api/v1/app-secrets-admin/revoke` credential-revoke endpoint.

└──▷ **TRY IT**

Revoke a credential after deleting a shared secret — now two separate calls since deletion no longer implies revocation.

```
$ curl -X DELETE
'https://<tenant>.conductor.one/api/v1/secrets/<vault_id>' \
  -H 'Authorization: Bearer <token>'

curl -X POST 'https://<tenant>.conductor.one/api/v1/app-secrets-
admin/revoke' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vault_id": "<vault_id>"}'
```

Diagnose a stuck provisioning task as a Super Admin without touching any live state.

📍 In the C1 console, open the stuck task's detail page, then select `Ask C1AI` and type: 'Why is this task stuck and what should I do next?'

Automatically assign access-review findings to the right owner using a Transformation rule instead of manual triage.

📍 In the C1 console, go to Automation › Transformation rules, open or create a rule, add a new step, select `Set assignee`, and specify the target user — this assignment will override any automatic assignee for matched findings.

Adds `POST /api/v1/app-secrets-admin/revoke` endpoint to queue credential revokes separately from secret deletion, replacing the previous combined behavior of `DELETE /api/v1/secrets/{vault_id}`.

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

Introduces `TYPE_SPEND_REMEDY = 7` as a new task action type; JS/TS clients must handle spend-control revision values as strings.

Adds Terraform resources for managing hosted and external MCP sources — including details, authentication, and tool enablement policy — as code via C1's Terraform provider.

Launches EU data residency instance, selectable at domain registration, affecting all tenant-specific URLs including login, redirect URIs, API and webhook endpoints, MCP server, and CLI configuration.

`+14 more`

└─▷ **BREAKING ON UPGRADE**

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

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud launched SaaS Management for shadow-IT discovery and app governance, added a Windows Device Patch Report with CSV export, extended AI Assistant with a Knowledge Base and AI Gateway with MCP server discovery, and introduced new OAuth client authentication options for agents — alongside breaking API changes that move user reactivation to a POST endpoint and rename Dynamic Groups filter operators.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 **OAuth client authentication types for agents** NEW 83

An agent's Authentication tab now offers Client secret basic and Client secret post client authentication type options, controlling how the agent platform sends the client ID and secret when requesting an OAuth token from JumpCloud. Client secret post is needed for platforms like Amazon Bedrock AgentCore that require credentials in the request body rather than the authorization header.

Match an agent's OAuth flow to Amazon Bedrock AgentCore, which requires credentials in the request body rather than the authorization header.

> 📍 In the Admin Portal, open the agent's Authentication tab, and under client authentication type select Client secret post so credentials are sent in the request body when requesting a token from JumpCloud.

Register an AI agent and configure it with OAuth credentials so it can access SSO applications through AI Gateway — without touching any human user accounts.

> 📍 1. Go to Identity Management › Agents and click + Add. 2. Enter the agent name, owner, and company email on the Details tab. 3. On the Authentication tab, generate OAuth client credentials; set client authentication type to Client secret post if your external platform (e.g. Amazon Bedrock AgentCore) sends credentials in the request body. 4. Under Identity Management › Agent Groups, create an Agent Group, add the agent, and assign SSO applications. 5. On the agent Applications tab, use Manage connections to authorize MCP servers on behalf of the agent.

— Exact UI path, field values, and example integration named.

Release Notes 2026 (2026-09-22) changelog-20260922-3d5122e5

02 **Device Patch Report for Windows** NEW 83

The Device Patch Report (Windows) shows per-device `Installed`, `Pending`, and `Failed` status for each Windows KB patch, with a preview of up to 500 rows in the portal and full CSV export of the complete result set.

Investigate patch compliance for a specific Windows device by previewing KB status in the portal, then downloading the full org dataset for reporting.

> 📍 In the Admin Portal, go to Device Patch Report (Windows), review the Installed, Pending, and Failed KB status per device (up to 500 rows in preview), then click Download to export the complete result set as CSV.

— Exact statuses, row limit, and export path given. `Release Notes 2026 (2026-09-22)`

---

03 **Global MFA Requirement for MSP admins** NEW 81

MSP admins can enforce MFA for all admins in the Multi-Tenant Portal via Account > Settings > `Global MFA Requirement` toggle, closing the gap where individual admin accounts might have MFA disabled.

Enforce Global MFA for all MSP admins in the Multi-Tenant Portal to close the gap where individual admin accounts might have MFA disabled.

> 📍 In the Admin Portal, go to Account › Settings and enable the `Global MFA Requirement` toggle under the `Global MFA Requirement` section.

— Exact toggle name and navigation path given. `Release Notes 2024 (2026-09-22)`

---

04 **MCP server discovery and AI Gateway registration** NEW 75

Under Access > AI & SaaS Management > Applications, a new MCP Connection column flags apps with a `Discovered` status; admins can review the discovered MCP servers and click Add MCP Server to register them in AI Gateway.

Check which managed applications have MCP servers discovered on devices, then connect one to AI Gateway — all without leaving the SaaS management view.

> 📍 1. Go to Access › AI & SaaS Management › Applications. 2. Locate the MCP Connection column; look for apps showing `Discovered`. 3. Click the `Discovered` status to review the discovered MCP servers. 4. Click Add MCP Server to register the server in AI Gateway. 5. To review details later, select the application and click View next to MCP Connection in the right-hand menu.

— Exact menu path and UI labels given, only from example. `Release Notes 2026 (2026-09-22)`

---

05 **Ad hoc command execution on devices** NEW 75

Admins can run an ad hoc shell command on a specific managed device via Devices > select device > Commands > Run Ad Hoc Command, without pre-binding it to a command object, and view output immediately in a results panel.

Run an ad hoc remediation command on a specific managed device without pre-binding it to a command object.

> 📍 In the Admin Portal, go to Devices › select a device › Commands › Run Ad Hoc Command, enter your shell command, and click Run to execute immediately and view output in the results panel.

— Exact menu path and behaviour given. Release Notes 2023 (2026-09-22)

---

06 **Preview dynamic group membership changes before applying** NEW

72

Dynamic Groups gain a 'Review membership changes before they take effect' option, letting admins preview membership updates — e.g. a group scoped by a Public IP Address rule — before they are applied.

Automatically keep a Remote Workers device group current without manual intervention by scoping membership to devices with a known public IP range.

> 📍 In the Admin Portal, go to Device Groups › Create Group › set Type to Dynamic, add a rule for Public IP Address matching your corporate egress range, then enable 'Review membership changes before they take effect' to preview before applying.

— Exact toggle name and example rule given, with steps. Release Notes 2023 (2026-09-22)

---

07 **Forced Re-authentication for SSO sessions** NEW 70

Supports Forced Re-authentication via `ForceAuthn='true'` in SAML and `prompt=login` in OIDC, prompting fresh credential entry even when an active session already exists.

— Exact protocol parameters named but no setup steps. Release Notes 2025 (2026-09-22)

---

08 **RSS feeds for agent release notes and bug fixes** NEW

63

Admins can subscribe separately to an Agent Release Notes RSS Feed and a Bug Fixes RSS Feed from the JumpCloud Release Notes and Bug Fixes page, to receive operational updates automatically in an RSS reader.

Subscribe to only agent release notes and bug fix reports to keep a dedicated RSS reader channel scoped to operational changes.

> 📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click Subscribe to Agent Release Notes RSS Feed to receive agent-specific updates. 3. Click

Subscribe to Bug Fixes RSS Feed to receive bug fix reports. 4. Open your RSS reader and add the resulting feed URLs to begin receiving notifications automatically.

— Exact feed names and subscription steps given. Subscribe to New Feature and Product Up…

THINNER COVERAGE BELOW

09 **Shared folder management and manual sync in Password Manager** NEW 52

Admins can manage shared folder assignments and permissions from the Admin Portal in JumpCloud Password Manager, and end users can manually sync individual folders from the desktop app.

— Describes capability and actor but no exact UI steps. Release Notes 2023 (2026-09-22)

10 **JumpCloud Go session auto-refresh** IMPROVED 43

JumpCloud Go sessions now automatically refresh on device login or unlock, removing additional verification prompts for JumpCloud-protected browser resources.

— Explains behaviour change but no configuration detail. Release Notes 2024 (2026-09-22)

11 **Premium Support Live Chat in Admin Portal** NEW 40

Premium Support Live Chat is now available directly in the Admin Portal, automatically recognizing Premium Support status without requiring a PIN.

— States what changed but no navigation or config detail. Release Notes 2022 (2026-09-22)

12 **macOS Dock and Lock Screen policy enhancements** IMPROVED 38

Adds macOS Dock Settings and Items policy updates, and enhancements to the Lock Screen (Screen Saver) policy for managed macOS devices.

— Names two policy areas but no specific settings changed.

Release Notes 2025 (2026-09-22)

13 **Multi-select and sorting in Password Manager desktop app** IMPROVED 25

Adds multi-select functionality and item sorting to the Password Manager desktop app.

— Names the change but no interaction detail. Release Notes 2024 (2026-09-22)

14 **Inactivity timer in Remote Assist** NEW 21

Adds an inactivity timer to the Remote Assist app.

— Bare one-line mention with no mechanism or config. Release Notes 2024 (2026-09-22)

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns 404; all scripts, automations, and integrations must switch to `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups API `memberQuery` filter operators `$le` and `$ge` are deprecated; API requests must use `$lte` and `$gte` respectively (Admin Portal-managed Dynamic Groups are not affected).

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

JumpCloud: Adds a required install success exit code field when deploying Windows EXE updates via the Private Repository, so JumpCloud can confirm installation succeeded.

└──▷ HOW TO FIND IT

Investigate patch compliance for a specific Windows device by previewing KB status in the portal, then downloading the full org dataset for reporting.

📍 In the Admin Portal, go to Device Patch Report (Windows), review the Installed, Pending, and Failed KB status per device (up to 500 rows in preview), then click Download to export the complete result set as CSV.

Match an agent's OAuth flow to Amazon Bedrock AgentCore, which requires credentials in the request body rather than the authorization header.

📍 In the Admin Portal, open the agent's Authentication tab, and under client authentication type select `Client secret post` so credentials are sent in the request body when requesting a token from JumpCloud.

Ground AI Assistant in org-specific runbooks and SOPs so it can answer policy questions with attributed sources.

📍 In the Admin Portal, navigate to AI Assistant, open the Knowledge Base section, and upload up to 25 internal documents (PDF, TXT, DOCX, or MD, max 10 MB each) — AI Assistant will cite the source document in any response that draws on uploaded content.

Adds Client secret basic and Client secret post client authentication type options on an agent's Authentication tab, controlling how the agent platform sends the client ID and secret when requesting an OAuth token from JumpCloud.

Release Notes 2022 (2026-09-22) BLOG ↗ STABLE

JumpCloud: Introduces a Linux Forbidden Services policy to remove unsupported service types and reduce security risk on Linux fleets.

└──▷ HOW TO FIND IT

Enforce Chrome security patches across every OS in your fleet with a single policy instead of per-OS rules.

📍 In the Admin Portal, go to Device Management › Policies › New Policy, select Browser Patch Management and choose the Universal (macOS, Windows, Linux) policy type, configure the update schedule, and save.

Harden Linux fleets by disabling risky or unnecessary system services across all managed Linux devices.

📍 In the Admin Portal, go to Device Management › Policies › New Policy, select Linux Forbidden Services, add the service names to block, assign the policy to the relevant Linux device group, and save.

Block macOS Ventura upgrades on managed devices until your org is ready to certify compatibility.

📍 In the Admin Portal, go to Device Management › Policies › New Policy, select the macOS software management policy to Delay or Block macOS Ventura, assign it to your macOS device groups, and save.

Adds Premium Support Live Chat directly in the Admin Portal, automatically recognizing Premium Support status without requiring a PIN.

Release Notes 2023 (2026-09-22) BLOG ↗ STABLE

JumpCloud: Enables admins to manage shared folder assignments and permissions from the Admin Portal in JumpCloud Password Manager, and lets users manually sync individual folders from the desktop app.

└──▷ HOW TO FIND IT

Automatically keep a Remote Workers device group current without manual intervention by scoping membership to devices with a known public IP range.

📍 In the Admin Portal, go to Device Groups › Create Group › set Type to Dynamic, add a rule for Public IP Address matching your corporate egress range, then enable

'Review membership changes before they take effect' to preview before applying.

Run an ad hoc remediation command on a specific managed device without pre-binding it to a command object.

> 📍 In the Admin Portal, go to Devices › select a device › Commands › Run Ad Hoc Command, enter your shell command, and click Run to execute immediately and view output in the results panel.

Adds billing visibility for MSP customers showing available and consumed user licenses and the ability to download their JumpCloud contract from the Multi-Tenant Portal.

Release Notes 2024 (2026-09-22) BLOG ↗ STABLE

JumpCloud: Adds manual application entry to SaaS Management so admins can track, warn, and block access to apps before automatic discovery runs.

└──▷ **TRY IT**

Reactivate a user account using the new POST endpoint after migrating away from the deprecated GET request.

```
$ curl -s -X POST
'https://console.jumpcloud.com/systemusers/reactivate/<id>' \
  -H 'x-api-key: <api_key>' \
  -H 'Content-Type: application/json' \
  -d '{"activation_email": "user@example.com"}'
```

Enforce Global MFA for all MSP admins in the Multi-Tenant Portal to close the gap where individual admin accounts might have MFA disabled.

> 📍 In the Admin Portal, go to Account › Settings and enable the `Global MFA Requirement` toggle under the `Global MFA Requirement` section.

Launches SaaS Management, providing discovery and tracking of SaaS app use via the JumpCloud Go browser extension, shadow-IT categorization, customizable warning/blocking messaging, and licensing information tracking.

Enables JumpCloud Go sessions to automatically refresh on device login or unlock, removing additional verification prompts for JumpCloud-protected browser resources.

Adds inactivity timer to Remote Assist app.

Adds multi-select functionality and item sorting to the Password Manager desktop app.

JumpCloud Go is no longer required at the organization level for SaaS Management app discovery; the JumpCloud Go browser extension alone suffices for end users.

└──▷ **BREAKING ON UPGRADE**

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns 404; all scripts, automations, and integrations must switch to `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups API `memberQuery` filter operators `$le` and `$ge` are deprecated; API requests must use `$lte` and `$gte` respectively (Admin Portal-managed Dynamic Groups are not affected).

`Release Notes 2025 (2026-09-22)` **BLOG ↗** STABLE

JumpCloud: Adds MTP support in the Customer Portal — MTP Admins can access, view, and create support cases across all managed tenant organizations.

Supports Forced Re-authentication via `ForceAuthn='true'` in SAML and `prompt=login` in OIDC, prompting fresh credential entry even when an active session exists.

Adds macOS Dock Settings and Items policy updates and Lock Screen (Screen Saver) policy enhancements for managed macOS devices.

`Release Notes 2026 (2026-09-22)` **BLOG ↗** STABLE

JumpCloud: Adds Client secret basic and Client secret post client authentication type options on an agent's Authentication tab, controlling how OAuth client credentials are sent when requesting tokens from JumpCloud.

└──▷ **HOW TO FIND IT**

Register an AI agent and configure it with OAuth credentials so it can access SSO applications through AI Gateway — without touching any human user accounts.

📍 1. Go to Identity Management › Agents and click + Add. 2. Enter the agent name, owner, and company email on the Details tab. 3. On the Authentication tab, generate OAuth client credentials; set client authentication type to Client secret post if your external platform (e.g. Amazon Bedrock AgentCore) sends credentials in the request body. 4. Under Identity Management › Agent Groups, create an Agent Group, add the agent, and assign SSO applications. 5. On the agent Applications tab, use Manage connections to authorize MCP servers on behalf of the agent.

Check which managed applications have MCP servers discovered on devices, then connect one to AI Gateway — all without leaving the SaaS management view.

📍 1. Go to Access › AI & SaaS Management › Applications. 2. Locate the MCP Connection column; look for apps showing `Discovered`. 3. Click the `Discovered` status to review the discovered MCP servers. 4. Click Add MCP Server to register the server in AI Gateway. 5. To review details later, select the application and click View next to MCP Connection in the right-hand menu.

Introduces the Device Patch Report (Windows) in JumpCloud Reports, providing per-device visibility into KB patch installation status (`Installed`, `Pending`, `Failed`) with preview of up to 500 rows and full CSV export.

Adds in-place version updating for custom applications in the Private Repository via `Update Version` on the application's Details tab, including EXE install success exit code configuration for Windows.

`What's New: Password Manager Desktop App (2026-09-22)` **BLOG ↗** **STABLE**

└──▷ **HOW TO FIND IT**

Activate Password Manager in the EU region to ensure credentials are stored and synced within EU infrastructure.

📍 On the desktop app activation page, select `EU` as your region before entering your activation code. For the browser extension, select `EU` on the pairing page before pairing.

└──▷ **BREAKING ON UPGRADE**

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

The easiest way to add authentication and user management to your application. Purpose-built for React, Next.js, Remix, and “The Modern Web”.

> Clerk's biggest ship this window is a full command-line workflow for agent-driven, zero-account Next.js setup, alongside GA of Client ID Metadata Documents with new admission controls and a dedicated SMS delivery logs view.

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

01 **Clerk CLI for agent-driven project setup** NEW 95

`npx -y clerk@latest init` provisions a claimable Clerk application, writes `NEXT_PUBLIC_CLERK_PUBLISHABLE_KEY` and `CLERK_SECRET_KEY` to `.env.local`, and configures the SDK, `<ClerkProvider>`, middleware, and auth routes with no Clerk account required; `init --app <application_id>` wires an existing application in instead of provisioning a new one. `apps list --json` enumerates existing applications by name and ID, `enable orgs` turns on organizations on a not-yet-claimed application, and `config patch` updates configuration on a not-yet-claimed application. The generated middleware file is named `proxy.ts` on Next.js 16+ and `middleware.ts` on Next.js 15 and below (identical contents), the CLI runs via `npx -y clerk@latest`, `pnpm dlx clerk@latest`, `bunx clerk@latest`, or `yarn dlx clerk@latest` (Yarn 2+), and it detects agent environments to run `init` non-interactively so no keys pass through the conversation.

Bootstrap Clerk in a new or existing Next.js project without a Clerk account — keys are written to `.env.local` automatically.

```
$ npx -y clerk@latest init
```

Connect a project to an existing Clerk application by choosing from all applications visible to the signed-in account.

```
$ npx -y clerk@latest apps list --json
# Review names and IDs, then:
npx -y clerk@latest init --app <application_id>
```

Enable organizations on a freshly provisioned, not-yet-claimed application before handing it to the user to claim.

```
$ npx -y clerk@latest enable orgs
```

— Full set of runnable CLI commands and flags named `product docs`

02 **Client ID Metadata Documents reach GA with admission controls**

NEW 75

Client ID Metadata Documents (CIMD) are now generally available for every Clerk application, letting MCP and other public OAuth clients connect with a URL-based identity without a support request. Operators can further restrict OAuth flows via a new `Client admission` control set to `Pre-registered clients only`, and enable CIMD in authorization server metadata via a `Publish CIMD support` toggle under `Client onboarding` on the OAuth applications settings page.

Lock your MCP server's OAuth flow so only clients you have explicitly reviewed can authenticate, using CIMD admission controls.

> 📍 In the Clerk Dashboard, go to OAuth applications settings › Client onboarding, enable `Publish CIMD support`, then set `Client admission` to `Pre-registered clients only` and add trusted client ID URLs from the Applications tab.

— Names exact settings and toggle but only UI steps given

`Client ID Metadata Documents are now ge…` `changelog-20260922-3af810e3`

03 **SMS delivery event logs** NEW 75

Application Logs now include `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` SMS delivery events, with a normalized `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) on failure events. A dedicated SMS logs view lets you filter by event type or destination phone number and see full per-message trace timelines.

Investigate why a user did not receive an SMS verification code by filtering the SMS logs to only failed delivery events for their number.

> 📍 In the Clerk Dashboard, go to Logs › SMS logs view, filter by event type `sms.failed`, then enter the destination phone number to pull up that recipient's full delivery trace.

— Named event types and reason values, UI-only path to use `changelog-20260922-3af810e3`

`changelog · Sep 22, 2026` CHANGELOG ↗ LATEST STABLE

Clerk adds Client ID Metadata Documents (GA) and SMS delivery event logs with per-event failure reasons

└──▷ **HOW TO FIND IT**

Investigate why a user did not receive an SMS verification code by filtering the SMS logs to only failed delivery events for their number.

> 📍 In the Clerk Dashboard, go to Logs › SMS logs view, filter by event type `sms.failed`, then enter the destination phone number to pull up that recipient's full delivery trace.

Lock your MCP server's OAuth flow so only clients you have explicitly reviewed can authenticate, using CIMD admission controls.

> 📍 In the Clerk Dashboard, go to OAuth applications settings › Client onboarding, enable `Publish CIMD support`, then set `Client admission` to `Pre-registered clients only` and add trusted client ID URLs from the Applications tab.

Adds `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` SMS delivery events to Application Logs, with a normalized `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) on failure events.

Adds a dedicated SMS logs view that lets you filter SMS delivery events by event type (e.g. `sms.failed`) or destination phone number, with full per-message trace timelines.

Adds `Client admission` control to CIMD settings, allowing operators to restrict OAuth flows to pre-reviewed clients by setting `Client admission` to `Pre-registered clients only`.

Adds `Publish CIMD support` toggle to the OAuth applications settings page under `Client onboarding`, enabling CIMD in the authorization server metadata.

`Client ID Metadata Documents are now generally available (2026-09-21)` **BLOG ↗**

STABLE

Client ID Metadata Documents (CIMD) are now generally available for all Clerk applications, enabling URL-based OAuth client identity.

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ announced · Sep 22, 2026

PRESS ↗ RANK

Secure, scalable, and highly available authentication and user management for any app.

> Okta expanded its AI Agent security suite with seven new capabilities spanning discovery, runtime enforcement, connection governance, and access reviews for AI agents.

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

01 **Kill Switch expanded to Agent Gateway** IMPROVED 56

Extends the existing Kill Switch to the Agent Gateway runtime layer, enabling revocation of every active token held by a compromised agent and termination of every session in flight at the moment of deactivation.

— Clear before/after mechanism but no steps to invoke it launch-20260922-e248582d

02 **Agent SSO via Cross App Access** NEW 52

Extends Cross App Access to all SSO customers, allowing non-expiring keys to be swapped for short-lived, identity-governed tokens to reduce user consent fatigue.

— Names the token model change but no setup instructions launch-20260922-e248582d

03 **Agent Gateway runtime enforcement** NEW 51

Agent Gateway is a runtime enforcement point that sits in the execution path between an agent and its tool calls, enforcing policy and logging every interaction.

— Explains mechanism but no config or API to act on launch-20260922-e248582d

04 **Agent-to-Agent Connections governance** NEW 44

Sets rules for which agents can call other agents, scopes what each agent can access, and captures the handoff in an auditable chain.

— Describes policy behavior but no interface or command given launch-20260922-e248582d

05 **Shadow AI Agent Discovery for Endpoints** NEW 38

Surfaces unmanaged AI agents running on employee devices so they can be discovered and registered before becoming governance blind spots.

— Names the feature and goal but no discovery mechanism detailed launch-20260922-e248582d

06 **Configuration Designer for agent connections** NEW 37

A visual interface for mapping agent-to-resource connections so that each handoff is governed, scoped, and auditable.

— Names a UI tool but gives no navigation path or specifics launch-20260922-e248582d

07 **Resource Access Certifications for agents** NEW 34

Periodically reviews agent resource connections over time to help prevent standing and excessive permissions.

— States purpose but no review cadence or configuration detail launch-20260922-e248582d

└──▷ **ALSO FROM THESE RELEASES**

announced · Sep 22, 2026 PRESS ↗

Okta for AI Agents adds shadow endpoint discovery, Agent Gateway runtime enforcement, agent-to-agent connection rules, and visual Configuration Designer.

Adds Shadow AI Agent Discovery for Endpoints, which surfaces unmanaged AI agents running on employee devices so they can be discovered and registered before becoming governance blind spots.

Adds Agent Gateway, a runtime enforcement point that sits in the execution path between an agent and its tool calls, enforcing policy and logging every interaction.

Adds Agent-to-Agent Connections, which sets rules for which agents can call other agents, scopes what each agent can access, and captures the handoff in an auditable chain.

Adds Configuration Designer, a visual interface for mapping agent-to-resource connections so that each handoff is governed, scoped, and auditable.

Expands the Kill Switch to the Agent Gateway runtime layer, enabling revocation of every active token held by a compromised agent and termination of every session in flight at the moment of deactivation.

+2 more

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Network, Ory Hydra, Ory Kratos, Ory Keto, Ory Oathkeeper, Ory Polis v26.3.15 release

NOTES ↗ RANK

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

> Ory Network now lets workspaces register named contacts for technical, business, security, and privacy roles.

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

01 **Workspace contact roles in Ory Network** NEW 48

Ory Network workspaces can register technical, business, security, and privacy contacts so Ory knows who to reach for each purpose. Contacts are set from the workspace Settings page in the Ory Network console.

Assign contacts for each role in a workspace so the right person is reached for security incidents or billing questions.

> 📍 In the Ory Network console, open your workspace, go to Settings, and fill in the technical, business, security, and privacy contact fields.

— Names roles and UI path but no API or config key

Ory Network, Ory Hydra, Ory Kratos, Ory…

Ory Network, Ory Hydra, Ory Kratos, Ory Keto, Ory Oathkeeper, Ory Polis v26.3.15 relea

NOTES ↗ LATEST STABLE

Ory Network gains workspace contact registration for technical, business, security, and privacy roles.

└──▷ **HOW TO FIND IT**

Assign contacts for each role in a workspace so the right person is reached for security incidents or billing questions.

> 📍 In the Ory Network console, open your workspace, go to Settings, and fill in the technical, business, security, and privacy contact fields.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

SandboxAQ AQtive Guard

Feedback ⓘ

changelog · Sep 22, 2026 CHANGELOG ↗ RANK

SandboxAQ AQtive Guard is a quantum-safe cryptography tool that protects against quantum computing threats to encryption.

> AQtive Guard added support for deploying its on-premises packages to RHEL 10.

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

01 **RHEL 10 support for on-premises packages** NEW 35

AQtive Guard on-premises packages now support RHEL 10 as a deployment target.

— Names the OS target but no install steps or version details.

changelog-20260922-256719a4

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

AQtive Guard on-premises packages now available for RHEL 10.

WAS THIS USEFUL? 👍 👎

Ubiq Security

Feedback ⓘ snapshot-20260922 NOTES ↗ RANK

Learn how Ubiq discovers sensitive data, maps identity-to-data access, and enforces protection policies across applications, databases, APIs, and AI.

> Ubiq expanded structured encryption dataset coverage across nearly every client library, shipped a beta Rust SDK, a Kong API gateway plugin, and Apache Airflow operators, and added identity-provider and SIEM depth to its access-governance dashboard.

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

01 **Structured encryption for integers, dates, tokens, and strings across client libraries**

NEW 93

The JavaScript library gains async structured encryption via `EncryptNumberAsync`, `DecryptNumberAsync`, and `encryptNumberForSearchAsync` for 32-bit (–99,999,999 to 99,999,999) and 64-bit (–9,999,999,999,999,999 to 9,999,999,999,999,999) integer datasets; `EncryptDateTimeAsync`, `DecryptDateTimeAsync`, and `encryptDateTimeForSearchAsync` for DateTime values (1653-02-10 to 2286-11-20 UTC); `EncryptDateAsync`, `DecryptDateAsync`, and `encryptDateForSearchAsync` for Date-only values; `EncryptAsync`, `DecryptAsync`, and `encryptForSearchAsync` for Token datasets of 64 characters (input up to 40 chars, base62 output) and 128 characters (input up to 80 chars, base62 output); and the same methods for generic string datasets with base32 or base64 encoding and configurable padding. Objects are built via the new `buildStructuredAsync` method on `CryptographyBuilder`. This dataset type support (`int32`, `int64`, `date`, `datetime`, `token` 64/128-byte, and generic strings with Base32/Base64/no encoding) was also extended across the C, Node.js, Python, PHP, .NET, and Go libraries.

Encrypt a 32-bit integer field (e.g. an account ID) so it remains searchable in a database without exposing plaintext values.

```text
const credentials =
ubiq.UbiqFactory.readCredentialsFromFile('./credentials',
'default');
const configuration = ubiq.UbiqFactory.defaultConfiguration();
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

const tweakFF1 = [];
const int32 = await
ubiqEncryptDecrypt.EncryptNumberAsync('int32', 9, tweakFF1);
const searchValues = await
ubiqEncryptDecrypt.encryptNumberForSearchAsync('int32', 9,
tweakFF1);
// use searchValues to query encrypted column
```

— Names every method, dataset type, and exact numeric ranges `snapshot-20260922`

02 **Airflow operators for structured encryption in DAGs** NEW

88

Adds `UbiqEncryptOperator` and `UbiqDecryptOperator` for Apache Airflow to protect data inside DAGs. Each accepts a `credentials` parameter (via ubiq.configCredentials()), a `data_task` parameter referencing a prior task's `task_id`, and a `dataset_names` list mapping column positions to Ubiq dataset names, with empty strings for columns left unencrypted.

Encrypt specific columns (SSN, Birthdate) pulled from BigQuery in a DAG, then decrypt them in a downstream task — leaving other columns untouched.

python

```python
from operators.ubiq_encrypt_operator import UbiqEncryptOperator
from operators.ubiq_decrypt_operator import UbiqDecryptOperator
import ubiq_security as ubiq

creds = ubiq.configCredentials(config_file='ubiq-credentials',
profile='default')

get_data = BigQueryGetDataOperator(
    task_id='get_data',
    dataset_id=DATASET_NAME,
    table_id=TABLE_NAME,
    max_results=10,
    selected_fields='name,email,ssn,birthday',
)

enc_task = UbiqEncryptOperator(
    task_id='enc_task',
    credentials=creds,
    data_task=get_data.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate'],
)

dec_task = UbiqDecryptOperator(
    task_id='ubiq_dec_task',
    credentials=creds,
    data_task=enc_task.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate'],
)

get_data >> enc_task >> dec_task
```

— Full runnable DAG example with all parameter names `snapshot-20260922`

---

03 **Custom metadata on usage records** NEW 75

Adds `addReportingUserDefinedMetadata` method accepting a JSON string (max 1,024 characters) to attach custom metadata to usage records on structured and conventional encryption objects, e.g. tagging records with an application or tenant identifier for audit correlation.

Attach custom audit metadata to usage records so SOC teams can correlate encryption activity with a specific application or tenant.

```text
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

ubiqEncryptDecrypt.addReportingUserDefinedMetadata('{"app":"payme
nts-service","tenant":"acme-corp"}');
```

— Named method, limit, and a concrete use case `product docs`

04 **Key number introspection in Go library** NEW 75

Adds `CipherAndKeyNumber`, `GetCurrentKeyNumber`, and `GetKeyNumber` functions to the Go library to return the key number behind a ciphertext without decrypting it, useful for key rotation auditing.

Inspect the key number behind a ciphertext in Go without decrypting, useful for key rotation auditing.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
```

— Three named functions plus a runnable call and use case `snapshot-20260922`

05 **Dataset cache prefetch in PHP library** NEW 72

Adds new `loadCache` function to the PHP library to prefetch dataset configuration and encryption keys for many datasets in a single call, reducing first-use latency.

Warm up all dataset keys at PHP application startup with `loadCache` to eliminate per-request key-fetch latency when encrypting across many datasets.

```php
$ubiq->loadCache($datasetNames);
```

— Named function with runnable one-line call and stated benefit `snapshot-20260922`

06 **Rust SDK (Beta)** NEW 68

Adds a new Ubiq Rust SDK (Beta), available on crates.io as `ubiq-rust`, supporting structured and unstructured encryption, integer and date/time datasets, `encrypt_for_search`, and key zeroization.

Use the new Rust SDK to encrypt a sensitive field with structured encryption in a Rust service — install the crate then call encrypt on a dataset.

```toml
[dependencies]
ubiq-rust = "1.0.0"
```

— Names crate, capabilities, but no full usage sample `snapshot-20260922`

07 **IDP-based credential creation** NEW 67

Adds `UbiqFactory.createCredentialsWithIdp` for creating credentials using an identity provider (IDP) via `<username>`, `<password>`, and `<host>`, enabling authentication through an SSO-enforced identity provider instead of static API keys.

Authenticate via an identity provider instead of static API keys, useful in SSO-enforced enterprise environments.

```text
const credentials =
ubiq.UbiqFactory.createCredentialsWithIdp(<username>, <password>,
<host>);
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)

.withConfigurationObject(ubiq.UbiqFactory.defaultConfiguration())
  .buildStructuredAsync();
```

— Named factory method with its parameters and use case `product docs`

THINNER COVERAGE BELOW

08 **Dashboard Integrations page for IDP and SIEM** IMPROVED 59

Updates the Dashboard Integrations page to configure IDP and SIEM connections in a single location, with support for self-signed IDPs, IDP group-to-Access Group mapping, and SCIM provisioning.

— Names a UI location and three specific capabilities `snapshot-20260922`

09 **Kong API gateway plugin for tokenization** NEW 52

Adds a new Ubiq plugin for the Kong API gateway providing field-level tokenization and detokenization of JSON request and response bodies with no application changes.

— Describes mechanism but no config example or endpoint `snapshot-20260922`

10 **Access Graph identity-to-dataset tracing** IMPROVED 50

Expands the Access Graph to include IDP groups and filter related nodes in both directions, enabling access tracing from either an identity or a dataset.

— Describes new traversal direction but no exact UI path `snapshot-20260922`

11 **IDP and SCIM attribution in Security History** IMPROVED 50

Adds IDP identity attribution to Security History events, records SCIM sync and attribute-mapping failures, and shows Access Group in event details.

— Names three distinct event-detail additions `snapshot-20260922`

12 **Grafana Loki SIEM integration** NEW 46

Adds Grafana Loki as a SIEM destination for Ubiq activity and usage events, available to Enterprise customers alongside existing Splunk support.

— Names the destination and audience but no setup detail `snapshot-20260922`

13 **Go input validation against dataset character set** IMPROVED 43

Adds improved input validation to the Go library — input is checked against the dataset character set before encryption, returning a clear error on mismatch.

— States mechanism but no error format or code sample `snapshot-20260922`

14 **.NET FF1 update to BITLEN** IMPROVED 41

Updates .NET structured encryption to use the `BITLEN` function in place of `LOG`, per NIST's updated FF1 recommendation.

— Names the function swap but no user-facing action `snapshot-20260922`

15 **Documented multithreading behavior for CryptographyBuilder** IMPROVED 41

Documents multithreaded-environment behavior: `CryptographyBuilder` objects are serialized and copied between Node.js worker threads rather than shared, requiring separate create/destroy cycles per thread.

— Explains behavior but is documentation-only, no example `product docs`

16 **Search and multi-select in Dataset Association dropdown** IMPROVED 40

Adds search and multi-select to the Dataset Association dropdown, and displays maximum input length for generic string datasets.

— Names UI widget but minimal behavioral detail `snapshot-20260922`

17 **New reference docs and integration guides** NEW 38

Adds new reference documentation for each structured dataset type (Date & Time, Date, Integer, Token, Formatted String, and Generic String) and new integration guides for Apache Airflow and Self-Signed IDP authentication.

— Lists doc topics but no content or code detail `snapshot-20260922`

`snapshot-20260922` **NOTES ↗** **LATEST** **STABLE**

Ubiq adds a Rust SDK, Kong gateway plugin, Airflow operators, Grafana Loki integration, and expanded dataset types across all client libraries.

└──▷ **USE IT**

Use the new Rust SDK to encrypt a sensitive field with structured encryption in a Rust service — install the crate then call encrypt on a dataset.

Cargo.toml

```
[dependencies]
ubiq-rust = "1.0.0"
```

Warm up all dataset keys at PHP application startup with `loadCache` to eliminate per-request key-fetch latency when encrypting across many datasets.

php

```
$ubiq->loadCache($datasetNames);
```

Inspect the key number behind a ciphertext in Go without decrypting, useful for key rotation auditing.

go

```
keyNum, err := ubiq.GetKeyNumber(ciphertext)
```

Adds new `loadCache` function to the PHP library to prefetch dataset configuration and encryption keys for many datasets in a single call, reducing first-use latency.

Adds new `CipherAndKeyNumber`, `GetCurrentKeyNumber`, and `GetKeyNumber` functions to the Go library to return the key number behind a ciphertext without decrypting it.

Adds new Ubiq Rust SDK (Beta) — available on crates.io as `ubiq-rust` — supporting structured and unstructured encryption, integer and date/time datasets, `encrypt_for_search`, and key zeroization.

Adds a new Ubiq plugin for the Kong API gateway providing field-level tokenization and detokenization of JSON request and response bodies with no application changes.

> Adds new Ubiq `encrypt` and `decrypt` operators for Apache Airflow to protect data inside Airflow DAGs.
>
> `+10 more`

WAS THIS USEFUL?  

Varonis

Feedback ⓘ Bring your own secret manager NOTES ↗ RANK

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

> Varonis expanded its integration footprint this window with new monitoring for GitHub Enterprise Cloud, Jira permission visibility, Confluence historical-content scanning, and self-service Snowflake PrivateLink configuration, alongside Private Cloud Data Collector improvements for secret management and Terraform deployment plus customizable on-demand scan prioritization.

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

01 Snowflake PrivateLink configuration in integration wizard

NEW 70

Adds support for configuring Snowflake integrations using AWS PrivateLink or Azure Private Link directly from the Varonis integration wizard, with automatic endpoint URL generation, pre-deployment connectivity validation, and remediation guidance.

— Names wizard and validation steps, exact UI path given

Configure Snowflake PrivateLink connect…

02 Private Cloud Data Collector secret management and Terraform deployment

NEW 62

Private Cloud Data Collectors can now be configured to use AWS Secrets Manager, Azure Key Vault, or `HashiCorp Vault` as external secret stores for collector credential storage, and can also be deployed to AWS ECS using Terraform for infrastructure-as-code management of collector environments.

— Names secret stores and deployment target but no exact config keys

Bring your own secret manager Deploy Private Cloud Data Collectors wi…

THINNER COVERAGE BELOW

03 Jira permissions visibility and overexposure detection NEW

55

Varonis tracks access to Jira content types — issues, comments, attachments, and worklogs — to surface who can access specific data and how that access is granted, and analyzes effective permissions across Jira to identify overexposed or broadly accessible resources.

— Names Jira content types but no UI or API path Identify overexposed Jira content

04 **GitHub Enterprise Cloud monitoring via GitHub App** NEW

50

Adds GitHub Enterprise Cloud monitoring via a GitHub App integration, scoping permissions to only what is required for monitoring and classification activities while providing enterprise-wide coverage and flexible deployment options.

— Names the integration mechanism but no setup steps

Monitor GitHub Enterprise Cloud via Git…

05 **Customizable on-demand scan prioritization** NEW 42

Adds on-demand scan prioritization, letting you select either a Prioritized Scan for immediate execution or a standard queued rescan, balancing urgency against resource efficiency.

— Explains the two options but gives no UI path or API

Customize on-demand scan prioritization

06 **Confluence page history scanning for sensitive data** NEW

40

Varonis now scans historical versions of Confluence pages to detect sensitive content, automatically tagging affected pages to support investigation and remediation workflows.

— Describes behavior but names no config or API surface

Investigate Confluence page history

Bring your own secret manager NOTES ↗ LATEST STABLE

Varonis Private Cloud Data Collectors now support AWS Secrets Manager, Azure Key Vault, and HashiCorp Vault for credential storage.

Customize on-demand scan prioritization NOTES ↗ STABLE

Varonis adds flexible on-demand scan prioritization, letting you choose between a Prioritized Scan and a standard queued rescan.

Investigate Confluence page history NOTES ↗ STABLE

Varonis now scans Confluence page history to surface sensitive data in older revisions and tag them for investigation.

Enables analysis of historical versions of Confluence pages to detect sensitive content, with automatic tagging of affected pages to support investigation and remediation workflows.

Deploy Private Cloud Data Collectors with Terraform NOTES ↗ STABLE

Varonis Private Cloud Data Collector can now be deployed to AWS ECS using Terraform.

`Identify overexposed Jira content` NOTES ↗ STABLE

Varonis adds Jira content visibility — track and analyze permissions on issues, comments, attachments, and worklogs.

Tracks access to Jira content types — issues, comments, attachments, and worklogs — to surface who can access specific data and how that access is granted.

Analyzes effective permissions across Jira to identify overexposed or broadly accessible resources.

`Monitor GitHub Enterprise Cloud via GitHub App` NOTES ↗ STABLE

Varonis can now monitor GitHub Enterprise Cloud via a GitHub App with least-privilege permissions and enterprise-wide coverage.

`Configure Snowflake PrivateLink connections directly in Varonis` NOTES ↗

STABLE

Varonis now supports self-service Snowflake PrivateLink configuration via the integration wizard for AWS and Azure.

WAS THIS USEFUL? 👍 👎

Bold Security

Feedback ⓘ announced · Sep 2, 2026 BLOG ↗

RANK

Bold Security runs a local AI agent on every endpoint that blocks insider threats, data leakage and shadow-AI use.

> Bold Security shipped a Wiz integration that feeds cloud resource exposure context into its endpoint agent's data-movement enforcement decisions.

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

01 Wiz integration for cloud-aware endpoint enforcement NEW

62

Bold syncs the Wiz cloud resource inventory — including repository, bucket, and database sensitivity and exposure levels — directly to the endpoint agent, and keeps this sync ongoing so allow/deny lists don't need to be maintained statically as cloud environments change. Enforcement decisions become exposure-aware, evaluating both the origin cloud resource and destination of a data transfer (e.g. distinguishing a move from an internal repo to an access-controlled bucket versus to a public S3 bucket), and enriched data lineage carries a file's source cloud resource (private repo, internal bucket, production database) as context from the moment data moves.

— Explains mechanism and scope well but no config/API names or UI path given.

launch-20260902-0521ae1a

> announced · Sep 2, 2026 BLOG ↗
>
> Bold Security integrates with Wiz to bring cloud resource exposure context into endpoint data-movement enforcement decisions.
>
> Integrates with Wiz to sync cloud resource inventory — including repository, bucket, and database sensitivity and exposure levels — directly to the Bold endpoint agent, so enforcement decisions reflect current cloud posture.
>
> Adds exposure-aware enforcement that evaluates both the origin cloud resource and the destination of a data transfer, distinguishing a move from an internal repo to an access-controlled bucket from the same move to a public S3 bucket.
>
> Adds enriched data lineage so a file's source cloud resource (private repo, internal bucket, production database) is carried as context to the endpoint from the moment data moves.
>
> Automatically syncs the Wiz cloud resource inventory on an ongoing basis, eliminating the need to maintain static allow/deny lists as cloud environments change.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ToolHive

Feedback ⓘ v0.51.0 NOTES ↗ RANK

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

> ToolHive v0.51.0 adds SPIFFE JWT-SVID authentication, letting MCP clients authenticate to MCP servers using workload identity.

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

01 **SPIFFE JWT-SVID authentication for MCP clients** NEW

35

ToolHive now authenticates MCP clients using SPIFFE JWT-SVIDs, enabling workload identity-based access control for MCP servers.

— Names the mechanism (SPIFFE JWT-SVID) but no config or command shown v0.51.0

v0.51.0 NOTES ↗ LATEST STABLE

ToolHive v0.51.0 adds SPIFFE JWT-SVID authentication for MCP clients.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.51.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.51.0
```

Authenticates MCP clients using SPIFFE JWT-SVIDs, enabling workload identity-based access control for MCP servers.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Sep 22, 2026 CHANGELOG ↗ RANK

The free and flexible app for your private thoughts.

> Obsidian 1.14.2 adds a command to open files from outside the vault, deeper Kanban folder-grouping support, and always-on fuzzy matching for link suggestions, while moving font-size sync out of `appearance.json`.

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

01 **Open file from outside the vault command** NEW 70

Adds the `Open file from outside the vault...` command to open any file on your computer in the current window (Desktop only).

Open a standalone Markdown report or log file that lives outside your vault for quick review without moving it.

> 📍 Open the command palette, search for 'Open file from outside the vault...', then navigate to and select any file on your filesystem.

— Named command with a usage example, but limited mechanism detail

changelog-20260922-e3fd1cda

02 **Always-on fuzzy matching for link suggestions** IMPROVED 65

The editor link suggestions (e.g. `[[`) and the quick switcher now always use fuzzy matching regardless of vault size, removing the 10,000-file threshold that previously triggered a simplified algorithm.

— Explains before/after threshold removal, no user action required

changelog-20260922-e3fd1cda

03 **Kanban folder-grouping and error handling** IMPROVED 60

Kanban boards grouped by `file.folder` now let cards be moved between folder columns, and creating a note in a column places it in the corresponding folder. Also adds error feedback when moving a card fails, such as when the note has invalid properties or the target folder is missing.

— Names grouping key and three behaviors, no config surface given

changelog-20260922-e3fd1cda

04 **Base font size no longer synced via Obsidian Sync**

BREAKING

60

Base font size is no longer saved to `appearance.json` and is now stored in device local storage, so it no longer syncs across devices via Obsidian Sync.

— Names config file and storage change but no migration guidance

changelog-20260922-e3fd1cda

THINNER COVERAGE BELOW

05 **Set Obsidian as default Markdown app** NEW 55

Obsidian now registers itself in the OS Open with menu and can be set as the default app for Markdown files (requires reinstalling the latest installer).

— Names OS integration and install requirement but no exact steps

changelog-20260922-e3fd1cda

06 **iOS 27 extra large widgets support** NEW 40

Adds support for the new extra large Quick Capture and View Note widgets on iOS 27.

— Names widget types but no setup instructions changelog-20260922-e3fd1cda

└──▷ BREAKING ON UPGRADE

! Base font size is removed from `appearance.json` and stored in device local storage — any synced font size preference via Obsidian Sync will no longer propagate across devices.

changelog · Sep 22, 2026 CHANGELOG ↗ LATEST STABLE

Obsidian 1.14.2 lets you open files from outside the vault and always uses fuzzy matching in large vaults.

└──▷ HOW TO FIND IT

Open a standalone Markdown report or log file that lives outside your vault for quick review without moving it.

📍 Open the command palette, search for 'Open file from outside the vault...', then navigate to and select any file on your filesystem.

Adds the `Open file from outside the vault...` command to open any file on your computer in the current window (Desktop only).

Obsidian now registers itself in the OS Open with menu and can be set as the default app for Markdown files (requires reinstalling the latest installer).

Kanban: Cards can now be moved between folder columns when grouped by `file.folder`.

Kanban: Creating a note in a column now places it in the corresponding folder when

grouped by `file.folder`.

Base font size is no longer saved to `appearance.json` and is now stored in device local storage, so it no longer syncs across devices via Obsidian Sync.

+3 more

└─▷ **BREAKING ON UPGRADE**

! Base font size is removed from `appearance.json` and stored in device local storage — any synced font size preference via Obsidian Sync will no longer propagate across devices.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 22, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue