THE CYBER TOOLCHAIN

NO. 089

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 29, 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 221 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 48 tools

▾ // THE BRIEF HIDE

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

Three unrelated things stand out. Knocknoc grew just-in-time network access from a handful of enforcement points to most of the firewall loud estate, so ports open per-session instead of standing open.

vim ? help

Qualys can now block a bad patch fleet-wide by rule and skip unreliable ones in zero-touch jobs. Snyk Container finally reports Go standard-library CVEs.

PROTECT

Open a port for one person for one session, on the firewall you already own

Enforcement now reaches DigitalOcean, Cloudflare, AWS WAF and EC2 Security Groups, Azure NSG, Sophos SFOS/XGS, F5 BigIP, Fortinet, PF, Windows Firewall and Linux IPSet from one Integrations hub, with PostgreSQL HA behind the server. The usual blocker for JIT network access is that it only covers one chokepoint and everything else stays allow-listed by IP forever; this covers enough of the estate to actually retire those rules, and the Trust Engine ships its decisions to SIEM.

Knocknoc

PROTECT

Stop a patch that breaks machines before it reaches the fleet

TruRisk Eliminate takes QQL-based rules that block specific patches environment-wide and uses reliability scoring to drop low-confidence patches out of zero-touch jobs, with curated interim mitigations for what stays unpatched. That is the missing half of automated patching: teams turn zero-touch off entirely after one bad Tuesday, and now they can exclude the offender instead. Iru adds per-device Windows update and patch compliance reporting so you can see who is actually current.

Qualys Patch Management · Iru

PROTECT

See Go stdlib CVEs in container scans instead of missing them

Vulnerabilities in the Go standard library are compiled into the binary and carry no manifest entry, so scanners routinely reported a clean image while the runtime itself was vulnerable. Snyk now synthesizes the stdlib as a dependency in results, which means the finding shows up in the existing triage queue rather than requiring a separate `govulncheck` pass someone remembers to run.

Snyk Container

IDENTIFY

Find the repos and cloud accounts your scanner never actually covered

> GitGuardian's Coverage Hub flags where integration, monitoring, scanning or ownership is missing across the monitored perimeter with one-click remediation, turning "we scan for secrets" into a list of the sources you demonstrably do not. Kubescape's multi-cluster fleet scanning rolls compliance into one report and adds scan-coverage metrics, so a clean posture score is distinguishable from a score computed over half the clusters.
>
> GitGuardian SaaS · kubescape

APPLICATION SECURITY

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ

Self-hosted runner version enforcement date has moved (2026-09-28) BLOG ↗

RANK

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

> GitHub shipped a REST API for tracking self-hosted runner version deprecations ahead of a tightened enforcement deadline, and made Claude Sonnet 5.5 generally available in Copilot with admin controls.

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

01 **Claude Sonnet 5.5 in GitHub Copilot** NEW 70

Claude Sonnet 5.5 is now generally available in GitHub Copilot across all supported clients and IDEs. Copilot Enterprise and Business admins can manage access via the model policy in Copilot settings, with new models automatically enabled under default model enablement unless the global default is off or the model is explicitly disabled.

Select Claude Sonnet 5.5 for an agentic coding task in VS Code to take advantage of its faster, lower-token execution on well-scoped feature work.

> 📍 In GitHub Copilot (VS Code, JetBrains, or github.com), open the model picker and select `Claude Sonnet 5.5`, then submit your Copilot Chat or coding agent prompt as normal.

As a Copilot Enterprise or Business admin, control which users can access Claude Sonnet 5.5 before the gradual rollout reaches your org.

> 📍 Navigate to your organization's Copilot settings › model policy, then enable or disable Claude Sonnet 5.5 under the default model enablement controls.

— Gives exact UI navigation and picker steps for use and admin control.

Claude Sonnet 5.5 in GitHub Copilot (20…

02 **REST API for runner version deprecations** BREAKING 60

A new REST API lets teams query registration and runtime deprecation dates for any self-hosted runner version, enabling automated alerts for fleets nearing their deadline. This accompanies a policy change: self-hosted runners below version 2.329.0 on GitHub Enterprise Cloud will not be able to register or reregister starting September 29, 2026,

and existing runners below the minimum runtime version will stop executing workflow jobs even if previously registered.

— Names deadline and version threshold but no endpoint path or method.

Self-hosted runner version enforcement …

└──▷ **BREAKING ON UPGRADE**

! Self-hosted runners below version 2.329.0 on GitHub Enterprise Cloud will not be able to register or reregister starting September 29, 2026; existing runners below the minimum runtime version will stop executing workflow jobs even if previously registered.

Self-hosted runner version enforcement date has moved (2026-09-28) **BLOG ↗**

LATEST **STABLE**

New REST API for runner version deprecations lets you query registration and runtime deprecation dates for any self-hosted runner version.

└──▷ **BREAKING ON UPGRADE**

! Self-hosted runners below version 2.329.0 on GitHub Enterprise Cloud will not be able to register or reregister starting September 29, 2026; existing runners below the minimum runtime version will stop executing workflow jobs even if previously registered.

Claude Sonnet 5.5 in GitHub Copilot (2026-09-28) **BLOG ↗** **STABLE**

Claude Sonnet 5.5 is now generally available in GitHub Copilot across all supported clients and IDEs.

└──▷ **HOW TO FIND IT**

Select Claude Sonnet 5.5 for an agentic coding task in VS Code to take advantage of its faster, lower-token execution on well-scoped feature work.

📍 In GitHub Copilot (VS Code, JetBrains, or github.com), open the model picker and select Claude Sonnet 5.5, then submit your Copilot Chat or coding agent prompt as normal.

As a Copilot Enterprise or Business admin, control which users can access Claude Sonnet 5.5 before the gradual rollout reaches your org.

📍 Navigate to your organization's Copilot settings › model policy, then enable or disable Claude Sonnet 5.5 under the default model enablement controls.

Copilot Enterprise and Business administrators can manage access to Claude Sonnet 5.5 via the model policy in Copilot settings, with new models automatically enabled

> under default model enablement unless the global default is off or the model is explicitly disabled.

WAS THIS USEFUL? 👍 👎

GitLab

Feedback ⓘ What's new in Git 2.56.0? (2026-09-28) BLOG ↗

RANK

The intelligent orchestration platform for DevSecOps, enabling teams and agents to ship trusted software at enterprise scale.

> Git 2.56.0 ships a batch of new subcommands and flags spanning history rewriting, branch cleanup, ref management, partial-clone maintenance, and object introspection.

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

01 New --delete-merged and --forked flags for git branch NEW

88

Adds `--delete-merged` flag to git-branch(1) to delete all local branches whose tips are reachable from their configured upstream, removing merged feature branches in one step. Adds `--forked` flag as a listing filter to show branches whose configured upstream matches a provided ref, remote name, or branch.

Clean up all local feature branches that have already been merged into origin/main, including branches forked from it rather than merged locally.

```
$ git branch --delete-merged origin
```

— Both flags named with mechanism and a runnable example.

What's new in Git 2.56.0? (2026-09-28)

02 Drop filtered blobs in git repack NEW

88

Adds `--drop-filtered` option to git-repack(1) to remove locally cached promisor blobs above the given filter threshold, recovering disk space in partial clones.

Recover disk space in a blobless partial clone by evicting large cached blobs so Git re-fetches them lazily only when needed.

```
$ git repack -a -d --filter=blob:limit=1k --drop-filtered
```

— Named flag with mechanism and a runnable example.

What's new in Git 2.56.0? (2026-09-28)

03 New git history drop subcommand NEW

80

Adds a `drop` subcommand to git-history(1) to remove a commit and rebase its descendants in a single command: `git history drop <commit>`.

Drop a specific commit from history and automatically rebase its descendants — no interactive rebase editor needed.

```
$ git history drop <commit>
```

— Names exact command and includes a runnable example.

What's new in Git 2.56.0? (2026-09-28)

---

04 **Remote object metadata lookup in git cat-file** NEW

78

Adds a `remote-object-info` command to git-cat-file(1)'s `--batch-command` mode to fetch object metadata (size, type) over Git protocol v2 without transferring the full object.

— Mechanism and protocol detail given, no usage example.

What's new in Git 2.56.0? (2026-09-28)

---

05 **New create, update, rename, delete subcommands for git refs**

NEW

73

Adds `create`, `update`, `rename`, and `delete` subcommands to git-refs(1), with `update` and `delete` accepting an optional expected old value for validation.

— All four subcommands named but no example given. What's new in Git 2.56.0? (2026-09-28)

---

06 **Linearize flag for git replay** NEW

70

Adds `--linearize` flag to git-replay(1) to enable rebasing commit ranges that contain merge commits by flattening the history.

— Flag and effect named, no example given. What's new in Git 2.56.0? (2026-09-28)

---

07 **Path fields added to git repo info** IMPROVED

60

Extends the `git repo info` subcommand with `path.gitdir` and `path.commondir` fields, each available in both absolute and relative forms.

— Fields named but no usage example provided. What's new in Git 2.56.0? (2026-09-28)

What's new in Git 2.56.0? (2026-09-28) BLOG ↗ LATEST STABLE

Git 2.56 adds `git history drop`, `git branch --delete-merged`, `git refs create/update/rename/delete`, `git repack --drop-filtered`, and more.

└──▷ TRY IT

Drop a specific commit from history and automatically rebase its descendants — no interactive rebase editor needed.

```
$ git history drop <commit>
```

Clean up all local feature branches that have already been merged into origin/main, including branches forked from it rather than merged locally.

```
$ git branch --delete-merged origin
```

Recover disk space in a blobless partial clone by evicting large cached blobs so Git re-fetches them lazily only when needed.

```
$ git repack -a -d --filter=blob:limit=1k --drop-filtered
```

Adds `drop` subcommand to git-history(1) to remove a commit and rebase its descendants in a single command: `git history drop <commit>`.

Adds `--delete-merged` flag to git-branch(1) to delete all local branches whose tips are reachable from their configured upstream, removing merged feature branches in one step.

Adds `--forked` flag to git-branch(1) as a listing filter to show branches whose configured upstream matches a provided ref, remote name, or branch.

Adds `create`, `update`, `rename`, and `delete` subcommands to git-refs(1), with `update` and `delete` accepting an optional expected old value for validation.

Adds `remote-object-info` command to git-cat-file(1)'s `--batch-command` mode to fetch object metadata (size, type) over Git protocol v2 without transferring the full object.

+3 more

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ Dev Machine Guard v1.17.0 (2026-09-29)

BLOG ↗ RANK

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

> StepSecurity's Dev Machine Guard v1.17.0 extends Secure Registry package config policies to PyPI/uv and Go, adds a fleet-wide Agent Plugins inventory for Claude Code and Codex, brings Insomnia into credential scanning, and redesigns the Compliance tab.

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

01 **Go Secure Registry package config policy** NEW 78

Adds Go package config policies that set the Secure Registry as the sole `GOPROXY` in the user's `go/env` file via the same shared `~/.netrc` credential, leaving `GOPRIVATE`, `GONOPROXY`, `GONOSUMDB`, and `GOSUMDB` untouched.

— Precise env vars and scope, but no CLI example given

Dev Machine Guard v1.17.0 (2026-09-29)

02 **PyPI/uv Secure Registry package config policy** NEW 75

Adds PyPI package config policies that point `pip` and `uv` at your tenant's Secure Registry by setting `index` in `pip.conf` and `uv.toml`, storing a shared credential in `~/.netrc`, and restoring the prior config when the policy is detached.

— Names exact config files and keys but no runnable command shown

Dev Machine Guard v1.17.0 (2026-09-29)

03 **Extra npm settings in package config policies** IMPROVED 75

Adds support for an optional list of extra `.npmrc` settings in npm package config policies, delivered to every assigned device inside the managed block, with tokens referenced as `${ENV_NAME}` from the developer's environment.

Roll out a scoped private registry for your org's npm packages alongside the Secure Registry, using an environment-variable token that StepSecurity never stores.

ini

```
@myorg:registry=https://npm.mycompany.internal
//npm.mycompany.internal/:_authToken=${MYORG_NPM_TOKEN}
```

— Includes a runnable .npmrc config example with token syntax

Dev Machine Guard v1.17.0 (2026-09-29)

04 **Agent Plugins inventory for Claude Code and Codex** NEW

70

Adds a new Agent Plugins inventory page under IDE & AI Agents that tracks Claude Code and Codex plugins fleet-wide, including marketplace/repository source, publisher, installed and enabled state per device, and the skills, commands, and MCP servers each plugin supplies.

Audit which Claude Code and Codex plugins are installed across your fleet before approving them for use in a compliance review.

> 📍 In the StepSecurity console, go to IDE & AI Agents › Agent Plugins to view each plugin's marketplace source, publisher, enabled state per device, and the MCP servers it supplies.

— Clear UI navigation given, no API or flag Dev Machine Guard v1.17.0 (2026-09-29)

05 **Insomnia credential scanning** NEW

65

Adds Insomnia to the credential inventory under a new API clients category, covering request authentication, credential headers, environment and secret variables, OAuth tokens, request history, and account/provider stores; secret variables in Insomnia's encrypted storage are reported as Protected.

— Thorough on scope but no action path given Dev Machine Guard v1.17.0 (2026-09-29)

06 **Redesigned Compliance tab per policy** IMPROVED

65

Redesigns the profile Compliance tab to show a compliance card per policy, one overall status per device, a device panel with per-policy status and reason, and highlights devices whose agent is too old for a policy.

— UI path implied and behaviour detailed, no exact navigation given

Dev Machine Guard v1.17.0 (2026-09-29)

07 **New coding tool and IDE detection** NEW

60

Adds detection for Grok Build, Kimi Code, Muse Code, Hermes Agent, and Oh My Pi as Coding Tools, and Kiro as an IDE, with their skills plus Kiro, Windsurf, Antigravity, and `OpenClaw` skills now appearing under Agent Skills.

— Many named tools but no usage or action described

Dev Machine Guard v1.17.0 (2026-09-29)

THINNER COVERAGE BELOW

08 **Agent Skills usage tracking** IMPROVED

45

Agent Skills now lists standalone Claude Code commands and shows a per-skill usage count drawn from Claude Code's recorded history.

— Describes what changed but no mechanism detail or path

Dev Machine Guard v1.17.0 (2026-09-29)

Dev Machine Guard v1.17.0 (2026-09-29) BLOG ↗ LATEST STABLE

Dev Machine Guard v1.17.0 adds PyPI, uv, and Go Secure Registry policies plus a Claude Code/Codex Agent Plugins inventory and Insomnia credential scanning.

└──▷ **HOW TO FIND IT**

Audit which Claude Code and Codex plugins are installed across your fleet before approving them for use in a compliance review.

📍 In the StepSecurity console, go to IDE & AI Agents › Agent Plugins to view each plugin's marketplace source, publisher, enabled state per device, and the MCP servers it supplies.

Roll out a scoped private registry for your org's npm packages alongside the Secure Registry, using an environment-variable token that StepSecurity never stores.

.npmrc

```
@myorg:registry=https://npm.mycompany.internal
//npm.mycompany.internal/:_authToken=${MYORG_NPM_TOKEN}
```

Adds PyPI package config policies that point `pip` and `uv` at your tenant's Secure Registry by setting `index` in `pip.conf` and `uv.toml`, storing a shared credential in `~/.netrc`, and restoring prior config when the policy is detached.

Adds Go package config policies that set the Secure Registry as the sole `GOPROXY` in the user's `go/env` file via the same shared `~/.netrc` credential, leaving `GOPRIVATE`, `GONOPROXY`, `GONOSUMDB`, and `GOSUMDB` untouched.

Adds support for an optional list of extra `.npmrc` settings in npm package config policies, delivered to every assigned device inside the managed block, with tokens referenced as `${ENV_NAME}` from the developer's environment.

Adds a new Agent Plugins inventory page under IDE & AI Agents that tracks Claude Code and Codex plugins fleet-wide, including marketplace/repository source, publisher, installed and enabled state per device, and the skills, commands, and MCP servers each plugin supplies.

Adds Insomnia to the credential inventory under a new `API clients` category, covering request authentication, credential headers, environment and secret variables, OAuth tokens, request history, and account/provider stores — secret variables in Insomnia's encrypted storage are reported as `Protected`.

> +3 more

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Training updates (2026-09-29) BLOG ↗ RANK

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

> Veracode's September 2026 training update centers on an upgraded Get Lesson Progress API, deeper Veracode Platform IAM integration for Security Labs team management, a wave of new OWASP Top 10:2025 and LLM security content, and several breaking changes including a new Security Labs URL and the retirement of the eLearning REST API and dozens of legacy courses.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 **Get Lesson Progress API enhancements** IMPROVED 78

Adds `updatedAfter` parameter to support incremental loads and `excludeNotStarted=true` to filter out lessons that have not been started, plus new `stack` and `updated_on` fields in the response; the API now returns `Not Started` when both `start_time` and `end_time` are null.

— Names exact API parameters and fields but no endpoint path or method.

Training updates (2026-09-29)

02 **Security Labs default URL changed** BREAKING 78

The default Security Labs URL has changed from `https://securitylabs.veracode.com/` to `https://securitylabs.veracode.com/security-labs/`; customers using SSO with a stand-alone account must update their Security Labs URL.

— Exact before/after URLs and a clear required migration step.

Training updates (2026-09-29)

03 **New Security Labs lesson content for OWASP 2025, LLM and containers**

NEW 63

Adds new OWASP Top 10:2025 Security Labs covering all 10 categories, including two new categories A03 Software Supply Chain Failures and A10 Mishandling of Exceptional Conditions; adds OWASP Top 10 for LLM applications lessons covering LLM01 through LLM09 in Flask; and adds Container Security in Docker lessons covering Users, Virtualization, Networking, Resources, Inventory, and Orchestration topics.

— Names all covered categories and topics but no way to start a lesson.

Training updates (2026-09-29)

04 **Deprecated eLearning courses and OWASP 2021 lessons removed**

DEPRECATED

63

33 deprecated eLearning courses (OWASP 2017, OWASP Top 10 (2021), Secure Coding for .NET, Java, JavaScript, PHP, Python — 6 courses each, Python 3 courses) are no longer available, with certification levels updated to include newer replacement courses; the OWASP Top 10 (2021) Security Labs lessons are also no longer available.

— Enumerates the removed courses precisely but no reader action path.

Training updates (2026-09-29)

THINNER COVERAGE BELOW

05 **Six new eLearning secure coding courses** NEW

55

Adds six new eLearning secure coding courses aligned to OWASP Top 10:2025 for .NET, Java, JavaScript, PHP, Python, and Security Awareness, available with English, Spanish, French, Portuguese, and Chinese captions.

— Lists languages and topics but no access instructions. Training updates (2026-09-29)

06 **Assessment quiz eLearning courses** NEW

55

Adds assessment quiz eLearning courses for User and application setup, Static Analysis, Dynamic Analysis, Scan and fix in your IDE, and SCA agent-based scanning, available in English, Portuguese, and Spanish.

— Names all course topics and languages but no access steps.

Training updates (2026-09-29)

07 **Team and IAM management in Security Labs** NEW

48

Organizations with more than 500 teams can now select which teams are available for use in Security Labs; Security Labs users and teams management is integrated with the Veracode Platform, with users authenticated through Veracode IAM; administrators and managers can also search teams within Security Labs.

— Describes what changed but not the underlying mechanism. Training updates (2026-09-29)

08 **Extended eLearning report date range** IMPROVED

45

Extends eLearning report date range selection to up to one year in the past.

— Gives concrete limit but no exact UI path. Training updates (2026-09-29)

09 **CWE tagging for Security Labs lessons** NEW

40

Adds CWE tagging for Security Labs lessons, enabling admins to search lessons by CWE and learners to see lessons linked to specific CWEs.

— Explains function but no search syntax or UI path. Training updates (2026-09-29)

10 **Security Labs terminology renamed** BREAKING

40

In Security Labs, `Teams` are now called `Organizations` and `Roles` are now called `Teams`; any documentation or automation referencing the old terminology must be updated.

— Names the terminology swap but no specifics on where it appears.

Training updates (2026-09-29)

---

11 **Veracode eLearning REST API end of support** DEPRECATED

30

The Veracode eLearning REST API has reached end of support and will no longer receive updates or technical support.

— States EOL status only, no migration path given. Training updates (2026-09-29)

---

12 **Security Labs community edition deprecated** DEPRECATED

30

The Security Labs community edition is deprecated as of July 31, 2026.

— Gives a firm date but no replacement or migration guidance.

Training updates (2026-09-29)

---

13 **Code syntax highlighting in lessons** NEW

25

Adds code syntax highlighting in Security Labs lessons tailored to programming languages.

— Single-line description with no further detail. Training updates (2026-09-29)

---

14 **In-app guidance for campaigns and reporting** NEW

20

Security Labs adds in-app guidance for creating campaigns and using reporting.

— Bare description with no mechanism or navigation detail. Training updates (2026-09-29)

└──▷ BREAKING ON UPGRADE

! The default Security Labs URL has changed from `https://securitylabs.veracode.com/` to `https://securitylabs.veracode.com/security-labs/`; customers using SSO with a stand-alone account must update their Security Labs URL.

! In Security Labs, `Teams` are now called `Organizations` and `Roles` are now called `Teams`; any documentation or automation referencing the old terminology must be updated.

! The Veracode eLearning REST API has reached end of support and will no longer receive updates or technical support.

! 33 deprecated eLearning courses (OWASP 2017, OWASP Top 10 (2021), Secure Coding for .NET, Java, JavaScript, PHP, Python — 6 courses each, Python 3 courses) are no longer available; certification levels have been updated to include newer replacement courses.

! The OWASP Top 10 (2021) Security Labs lessons are no longer available.

! The Security Labs community edition is deprecated as of July 31, 2026.

Training updates (2026-09-29) BLOG ↗ LATEST STABLE

Veracode Security Labs gains team selection for large orgs, Veracode Platform IAM integration, and updated Get Lesson Progress API with new parameters.

Adds `updatedAfter` parameter to the Get Lesson Progress API to support incremental loads.

Adds `excludeNotStarted=true` parameter to the Get Lesson Progress API to filter out lessons that have not been started.

Adds `stack` field and `updated_on` field to the Get Lesson Progress API response; shows `Not Started` when both `start_time` and `end_time` are null.

Enables organizations with more than 500 teams to select which teams are available for use in Security Labs.

Integrates Security Labs users and teams management with the Veracode Platform, authenticating Security Labs users through Veracode IAM (Identity and Access Management).

+10 more

└──▷ **BREAKING ON UPGRADE**

- ! The default Security Labs URL has changed from `https://securitylabs.veracode.com/` to `https://securitylabs.veracode.com/security-labs/`; customers using SSO with a stand-alone account must update their Security Labs URL.
- ! In Security Labs, `Teams` are now called `Organizations` and `Roles` are now called `Teams`; any documentation or automation referencing the old terminology must be updated.
- ! The Veracode eLearning REST API has reached end of support and will no longer receive updates or technical support.
- ! 33 deprecated eLearning courses (OWASP 2017, OWASP Top 10 (2021), Secure Coding for .NET, Java, JavaScript, PHP, Python — 6 courses each, Python 3 courses) are no longer available; certification levels have been updated to include newer replacement courses.
- ! The OWASP Top 10 (2021) Security Labs lessons are no longer available.
- ! The Security Labs community edition is deprecated as of July 31, 2026.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ

Spot your uncovered assets. Close your blindspot in one click! (2026-09-28)

BLOG ↗ RANK

GitGuardian SaaS is a cloud-based security platform that detects and prevents secrets leakage in git repositories and source code.

> GitGuardian's biggest addition this window is the Coverage Hub on the Monitored Perimeter page, which surfaces integration, monitoring, scanning and ownership gaps with one-click remediation; the perimeter page also gained saved views, column customization, team-ownership tracking, improved bulk actions, and automatic team-perimeter eviction rules for deleted, archived or unmonitored sources, alongside a breaking replacement of the Health filter with Has open incidents.

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

01 **Coverage Hub on the perimeter page** NEW 88

Adds the Coverage Hub to the Monitored Perimeter page: four summary cards (Integrations, Real time monitoring, Historical scans, Team ownership) quantify coverage gaps — installed integrations, share of sources monitored in real time, share scanned successfully, and share owned by a team — and link directly to filtered source lists. Card actions filter the sources table to sources needing action (unmonitored, unscanned, or unassigned) in one click, and a red badge counts sources GitGuardian cannot reach, with a Show more details view listing integrations and flagging unreachable ones.

Surface integrations with unreachable sources so you can restore monitoring coverage before secrets go undetected.

> 📍 In the perimeter page, look for the red badge on the Coverage Hub, then click Show more details to see which integrations have unreachable sources and click the card action to filter the table to those sources.

Identify all sources not yet assigned to a team and bulk-assign them to close ownership gaps.

> 📍 In the perimeter page, click the Owned by a team card action to filter the sources table to unassigned sources, select all, then choose a bulk action to assign

them. Alternatively, use the Team filter and select `No team assigned` to reach the same view.

— Names all four cards and gives exact click-through steps.

`Spot your uncovered assets. Close your …`

02 **Bulk actions on perimeter sources** IMPROVED 80

Adds a bulk actions bar for acting on many perimeter sources at once — `Monitor`, `Unmonitor`, and `Launch historical scan` — requiring confirmation when selecting more than 30 sources or when some sources will be skipped. Later grays out inapplicable actions with explanatory tooltips and lists skipped sources with reasons when monitoring, unmonitoring, or scanning more than 30 sources.

Kick off historical scans for every source that has never been scanned, confirming the batch when GitGuardian prompts on selections over 30 sources.

📍 In the Monitored Perimeter page, click `See unscanned sources` in the Historical scans card, use 'Select all N sources that match filters', choose the `Launch historical scan` bulk action, and confirm the dialog.

— Names each bulk action and a worked example with the 30-source threshold.

`Spot your uncovered assets. Close your …`

03 **Automatic team-perimeter eviction rules** NEW 75

Adds configurable automatic team-perimeter eviction rules — `Deleted sources`, `Archived sources`, and `Unmonitored sources` — so members lose access to a source and its incidents as soon as it leaves the team perimeter under one of these statuses. Managers toggle the `team perimeter setting` per status to keep or remove unmonitored and archived/deleted sources from team perimeters.

Automatically remove unmonitored sources from a team perimeter so stale sources do not clutter a team's incident view — while keeping manager access intact.

📍 In the GitGuardian console, go to Workspace Settings › Team perimeter and enable the `Unmonitored sources` setting so that sources which are no longer monitored are automatically removed from all team perimeters.

— Names each status toggle and the console path to enable it. `product docs`

04 **Saved views and column customization on perimeter page** NEW 70

Adds `saved views` to the perimeter page, including built-in non-editable views `Monitored`, `Critical sources`, `With open secret incidents`, `Scanning issues`, and `Without`

honeytoken ; a Columns control to show, hide, or reorder perimeter table columns; a Teams column (shown by default in the Monitored view) with a No team assigned filter value to identify unowned sources; and a Last scan column with a Skipped badge and hover explanation for scans that did not run.

.

— Enumerates every named view and column but no runnable step.

Spot your uncovered assets. Close your …

THINNER COVERAGE BELOW

05 **Health filter replaced by Has open incidents** BREAKING

50

Replaces the Health filter with Has open incidents on the perimeter page, with saved views using the old filter migrated automatically; the sources table is now sorted by open incidents by default, prioritising sources at risk.

— States the before/after but gives no migration steps.

Spot your uncovered assets. Close your …

└──▷ BREAKING ON UPGRADE

! The Health filter on the perimeter page is replaced by Has open incidents ; saved views using the old filter are migrated automatically.

Spot your uncovered assets. Close your blindspot in one click! (2026-09-28)

BLOG ↗ STABLE

GitGuardian's new Coverage Hub surfaces unmonitored, unscanned, and unowned sources on the perimeter page with one-click remediation.

└──▷ HOW TO FIND IT

Identify all sources not yet assigned to a team and bulk-assign them to close ownership gaps.

📍 In the perimeter page, click the Owned by a team card action to filter the sources table to unassigned sources, select all, then choose a bulk action to assign them. Alternatively, use the Team filter and select No team assigned to reach the same view.

Surface integrations with unreachable sources so you can restore monitoring coverage before secrets go undetected.

📍 In the perimeter page, look for the red badge on the Coverage Hub, then click Show more details to see which integrations have unreachable sources and click the card action to filter the table to those sources.

Adds the Coverage Hub to the perimeter page: four cards show installed integrations, share of sources monitored in real time, share scanned successfully, and share owned by a team.

Adds card actions that filter the sources table to sources needing action — unmonitored, unscanned, or unassigned — in one click.

Adds a red badge on the Coverage Hub to count sources GitGuardian cannot reach, with a `Show more details` view listing integrations and flagging unreachable ones.

Adds a `Teams` column to the perimeter sources table showing the teams assigned to each source.

Adds a `No team assigned` option to the Team filter to surface sources without an owner.

`+3 more`

└─▷ **BREAKING ON UPGRADE**

! The `Health` filter on the perimeter page is replaced by `Has open incidents`; saved views using the old filter are migrated automatically.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

iLEAPP

Feedback ⓘ v2026.4.3 NOTES ↗ CODE ↗ RANK

iOS Logs, Events, And Plist Parser

> iLEAPP v2026.4.3 unifies GUI and CLI into a single `ileapp` binary, greatly expands the forensic image formats it can parse directly, and adds nine new iOS app parsers plus several new artifact types.

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

01 Raw image and filesystem support in `-t raw` IMPROVED

90

Extends `-t raw` to parse E01 and Ex01, SMART, AFF and AFD forensic images, Apple disk images (`.dmg` , `.sparseimage` , sparse bundles and segmented sets), encrypted Apple disk images, FTK Imager AD-encrypted images, and encrypted APFS volumes opened with a supplied password, without mounting or admin rights. Also adds NTFS alternate data stream reading under `-t raw` .

Parse an EnCase E01 acquisition directly — no mounting or admin rights needed — useful when working from a forensic image on a case workstation.

```
$ ileapp -t raw -i /evidence/acquisition.E01 -o /cases/output/
```

— Lists every supported format and includes a runnable command v2026.4.3

02 Unified `ileapp` binary across platforms BREAKING

80

iLEAPP now ships as a single `ileapp` binary for both GUI and CLI on all platforms — running without arguments launches the GUI, running with arguments runs the CLI — replacing the separate `ileappGUI` binary. Windows downloads are `-windows-x64-setup.exe` / `-windows-x64-portable.zip` and `-windows-arm64-setup.exe` / `-windows-arm64-portable.zip` ; macOS gets signed and notarised `-macos-arm64.dmg` / `-macos-x64.dmg` ; Linux gets `-linux-x64.AppImage` / `-linux-arm64.AppImage` . Automated tooling that invoked `ileappGUI` or the old file-naming scheme must be updated.

— Names exact new download files but no migration command v2026.4.3

03 Nine new iOS app parsers NEW

60

Adds iOS artifact parsers for Roblox, KakaoTalk (with message, attachment and phone field decryption and inline media), WeChat, Telekom Mail, BOTIM, VK Messenger,

Odnoklassniki, and SimpleX Chat (SQLCipher database decrypted with the keychain key).

— Names every app but no usage detail beyond that v2026.4.3

THINNER COVERAGE BELOW

04 **GUI support for .tar.xz archives** IMPROVED 50

Adds GUI support for accepting `.tar.xz` archives, matching the existing CLI support via `-t tar`.

— Names format and flag but no full command shown v2026.4.3

05 **Sysdiagnose Process tree artifact** NEW 45

Adds a Sysdiagnose Process tree artifact built from `ps.txt`.

— Names the source file but no usage instructions v2026.4.3

06 **Intelligence Platform artifact family for iOS** NEW 35

Adds a new Intelligence Platform artifact family for iOS covering the knowledge graph, expired interactions, behaviors timeline, and event locations.

— Names the four sub-artifacts but no mechanism v2026.4.3

07 **Snapchatter Records rename and friends-list parsing** IMPROVED 35

Renames the Snapchat Friends artifact to Snapchatter Records and adds parsing of the friends list.

— Clear rename and addition, no further mechanism given v2026.4.3

08 **Call History columns for iOS 26** NEW 30

Adds Call History columns for iOS 26.

— Bare mention of new columns, no specifics v2026.4.3

09 **qnxprobe dependency updated to 1.51** IMPROVED 30

Updates the bundled `qnxprobe` tool to version 1.51.

— Version bump only, no behaviour change described v2026.4.3

10 **iOS Health walking and running distance artifact** NEW 25

Adds an iOS Health artifact for walking and running distance.

— Thin one-line addition with no further detail v2026.4.3

11 **Biome sync merged events artifact** NEW 25

Adds a Biome sync merged events artifact.

— Thin one-line addition with no further detail v2026.4.3

└──▷ **BREAKING ON UPGRADE**

! The separate `ileappGUI` binary no longer exists; both GUI and CLI are now the single `ileapp` binary (GUI when run without arguments, CLI when run with arguments). Automated tooling that invokes `ileappGUI` must be updated.

! Download file names changed: the previous naming scheme (e.g. `ileappGUI-v*-Windows_x86_64.zip`) is replaced by new names such as `-windows-x64-setup.exe`, `-windows-x64-portable.zip`, `-macos-arm64.dmg`, `-linux-x64.AppImage`, etc.

v2026.4.3 NOTES ↗ CODE ↗ LATEST STABLE

iLEAPP v2026.4.3 adds unified `ileapp` binary for all platforms, E01/AFD/APFS image support via `-t raw`, and nine new iOS app parsers.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2026.4.3
https://github.com/abrignoni/iLEAPP.git
# already have the repo? check out this version:
$ git checkout v2026.4.3
```

└──▷ **TRY IT**

Parse an EnCase E01 acquisition directly — no mounting or admin rights needed — useful when working from a forensic image on a case workstation.

```
$ ileapp -t raw -i /evidence/acquisition.E01 -o /cases/output/
```

Extends `-t raw` to support E01 and Ex01, SMART, AFF and AFD images, Apple disk images (`.dmg`, `.sparseimage`, sparse bundles and segmented sets), encrypted Apple disk images, FTK Imager AD-encrypted images, and encrypted APFS volumes (opened with a supplied password).

Adds NTFS alternate data stream reading under `-t raw`.

Adds iOS artifact support for Roblox, KakaoTalk (with message/attachment/phone field decryption and inline media), WeChat, Telekom Mail, BOTIM, VK Messenger, Odnoklassniki, and SimpleX Chat (SQLCipher, keychain key).

Adds iOS Health walking and running distance artifact.

Adds Intelligence Platform artifact family for iOS: knowledge graph, expired interactions, behaviors timeline, and event locations.

+7 more

└─▷ **BREAKING ON UPGRADE**

! The separate `ileappGUI` binary no longer exists; both GUI and CLI are now the single `ileapp` binary (GUI when run without arguments, CLI when run with arguments). Automated tooling that invokes `ileappGUI` must be updated.

! Download file names changed: the previous naming scheme (e.g. `ileappGUI-v*-Windows_x86_64.zip`) is replaced by new names such as `-windows-x64-setup.exe`, `-windows-x64-portable.zip`, `-macos-arm64.dmg`, `-linux-x64.AppImage`, etc.

WAS THIS USEFUL?  

Mobile Verification Toolkit

Feedback ⓘ v2026.9.28

NOTES ↗ CODE ↗ RANK

MVT (Mobile Verification Toolkit) helps with conducting forensics of mobile devices in order to find signs of a potential compromise.

> MVT improved its bugreport forensics pipeline by making `check-bugreport` recover data from nested OEM archives and fixing a parsing bug that dropped system properties, and it added a new alert for suspicious accessibility service counts.

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

01 **Bugreport archive handling improvements** IMPROVED 65

`check-bugreport` now descends into OEM wrapper `.zip` archives when the outer archive lacks a dumpstate entry point, automatically finding and using the nested bugreport. A separate fix corrects cases where a section timing line was misidentified as a section boundary, restoring full `SYSTEM PROPERTIES` output on affected archives.

— Names the subcommand and fixed output section, no exact CLI flags given v2026.9.28

02 **Accessibility service count alert in dumpsys analysis** NEW 65

Dumpsys accessibility analysis now surfaces accessibility service counts (`installedServiceCount=N`) per user, raising a LOW alert when the number of services listed in the count does not match the number actually named in the dump.

— Names the field and alert severity but no example output shown v2026.9.28

v2026.9.28 NOTES ↗ CODE ↗ LATEST STABLE

MVT's `check-bugreport` now descends into OEM wrapper archives and recovers previously truncated dumpstate sections

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2026.9.28 https://github.com/mvt-project/mvt.git
# already have the repo? check out this version:
$ git checkout v2026.9.28
```

Extends `check-bugreport` to descend into OEM wrapper `.zip` archives when the outer archive lacks a dumpstate entry point, automatically finding and using the

nested bugreport.

Recovers properties lost when a section timing line was misidentified as a section boundary in bugreport parsing, restoring full `SYSTEM PROPERTIES` output on affected archives.

Surfaces accessibility service counts (`installedServiceCount=N`) per user in dumpsys accessibility analysis, raising a LOW alert when services are listed in the count but not named in the dump.

WAS THIS USEFUL? 👍 👎

macos-unifiedlogs

Feedback ⓘ v0.7.0 NOTES ↗ RANK

A cross platform parser for Apple UnifiedLogs!

Adds `lzbitmap` decompression and richer Simpledump/Statedump field extraction.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.7.0 https://github.com/mandiant/macos-UnifiedLogs.git
# already have the repo? check out this version:
$ git checkout v0.7.0
```

Adds `lzbitmap` decompression support for parsing log entries compressed with that algorithm.

Includes additional details extracted from Simpledump and Statedump log entries.

NEW SUPPORTED PLATFORM GoldenGate and iOS 27 Unified Log support

WAS THIS USEFUL? 👍 👎

cortex-analyzers

Feedback ⓘ 3.10.0 NOTES ↗ RANK

Cortex Analyzers Repository

v3.10.0 adds a ThreatFox analyzer, an Elastic Security alert-sync responder, and CheckPoint Harmony Email & Collaboration responders.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.10.0 https://github.com/TheHive-Project/Cortex-Analyzers.git
# already have the repo? check out this version:
$ git checkout 3.10.0
```

NEW INTEGRATION Adds a new ThreatFox analyzer for querying the Abuse.ch ThreatFox threat intelligence platform.

NEW INTEGRATION Adds an Elastic Security responder to sync alert statuses between Cortex and Elastic Security.

NEW INTEGRATION Adds CheckPoint Harmony Email & Collaboration responders for incident response actions against that platform.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

September 28, 2026 - Project name checks, Finding VEX retention, and faster Component re

BLOG ↗ RANK

Finite State Platform is a software supply chain security tool that analyzes firmware and software composition to identify vulnerabilities and compliance risks.

> Finite State Platform's latest release adds an audit trail for user deletions, speeds up Component replace/revert on large Versions, preserves VEX and triage data through reassessment, and cleans up Findings CSV export headers.

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

01 **USER_DELETED audit event on user deletion** NEW 90

When `DELETE /public/v0/users/{id}` succeeds, the platform now logs a `USER_DELETED` event to `GET /public/v0/audit` including the deleted user's `id` and `email`, and terminates that user's active sessions.

Verify that deleting a user produces an audit trail — confirm the `USER_DELETED` event and the user's id and email appear in the audit log.

```
$ curl -X DELETE
'https://platform.finitestate.io/public/v0/users/{id}' -H
'Authorization: Bearer <token>'

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

— Named endpoints, event type, and fields with a runnable verification example

September 28, 2026 - Project name check…

THINNER COVERAGE BELOW

02 **Faster Component replace and revert on large Versions**

IMPROVED 50

Component replace and revert operations on large Versions are now faster because license Policy is evaluated once per Component rather than once per finding.

— Explains mechanism but gives no command or config to act on

September 28, 2026 - Project name check…

03 **VEX and comment retention through reassessment** IMPROVED

50

Findings that carry a VEX status or a comment — including SBOM-supplied VEX and triage — are now preserved through reassessment triggered by exclude, include, replace, or revert operations.

— Clear behavior described but no interface or command to trigger it

September 28, 2026 - Project name check…

04 **Consistent column headers in Findings CSV exports** IMPROVED

35

Filtered, searched, and row-selected Findings CSV exports now use the same column header labels as the unfiltered export, instead of raw internal key names.

— States before/after but no export path or command given

September 28, 2026 - Project name check…

September 28, 2026 - Project name checks, Finding VEX retention, and faster Component

BLOG ↗ LATEST STABLE

Finite State Platform now logs `USER_DELETED` audit events, retains VEX/comments through reassessment, and speeds up Component replace on large Versions.

└──▷ **TRY IT**

Verify that deleting a user produces an audit trail — confirm the `USER_DELETED` event and the user's id and email appear in the audit log.

```
$ curl -X DELETE
'https://platform.finitestate.io/public/v0/users/{id}' -H
'Authorization: Bearer <token>'

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

Adds `USER_DELETED` audit event to `GET /public/v0/audit` — including the deleted user's `id` and `email` — when `DELETE /public/v0/users/{id}` succeeds, and terminates the deleted user's active sessions.

Speeds up Component replace and revert on large Versions by evaluating license Policy once per Component rather than per finding.

Preserves Findings that carry a VEX status or a comment (including SBOM-supplied VEX and triage) through reassessment after exclude, include, replace, or revert operations.

Filtered, searched, and row-selected Findings CSV exports now use the same column header labels as the unfiltered export instead of raw internal key names.

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

YARA-X

Feedback ⓘ v1.21.0 NOTES ↗ RANK

A rewrite of YARA in Rust.

Adds `Scanner.set_module_output` to Python bindings, a new `debug` command for atoms, and stdin-based scan list support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.21.0 https://github.com/VirusTotal/yara-x.git
# already have the repo? check out this version:
$ git checkout v1.21.0
```

Adds `Scanner.set_module_output` method to the Python bindings for controlling module output during scans.

Adds `debug` command for showing atoms extracted from YARA-X rules.

Supports reading scan target lists from stdin, enabling pipeline-driven scanning workflows.

NEW SUPPORTED PLATFORM Additional ELF machine types in the `elf` module

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

SO-CRATES

Feedback ⓘ v4.3.0 NOTES ↗ RANK

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

> SO-CRATES v4.3.0 bundles a fully offline CyberChef instance directly into the container and wires it into the UI at multiple entry points — stream payloads, Suricata-extracted files, and ad-hoc transcript selections — for air-gapped decode-and-pivot workflows.

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

01 **Send stream payload bytes to CyberChef** NEW 80

Adds a Send to CyberChef action in a stream's Payload panel that sends the exact bytes in Both directions, or just the Source or Dest side, preserving XOR keys, compressed data, and shellcode intact for analysis in the bundled CyberChef.

Analyse a full bidirectional TCP payload — shellcode, compressed data, XOR-keyed content — in an air-gapped CyberChef without it appearing in browser history.

📍 In the stream detail view, open the Payload panel, then click Send to CyberChef and choose Both, Source, or Dest to load the exact bytes into the bundled CyberChef instance.

— Names exact panel, direction options, and byte-level guarantees with a runnable UI flow v4.3.0

02 **Pivot-menu text selection with auto Magic decode** IMPROVED 65

Selecting any text in an ASCII transcript — base64 blobs, headers, or partial values — now surfaces a pivot menu with CyberChef (and other lookups) access. CyberChef's Magic operation is applied automatically to small payloads sent this way, including values sent from the pivot menu, to suggest decodings.

Decode a suspicious base64 blob or header value spotted in an ASCII transcript without leaving the SO-CRATES UI.

📍 Select the text (e.g. a base64 blob) in the ASCII transcript, then choose CyberChef from the pivot menu that appears — CyberChef opens with the selected text and, for small values, applies the Magic operation automatically.

— Describes trigger and auto-decode mechanism with a concrete usage walkthrough v4.3.0

THINNER COVERAGE BELOW

03 **Bundled, air-gapped CyberChef instance** NEW 55

CyberChef is now bundled into the container image (~48 MB added) so pivot-menu CyberChef lookups work with no internet access. The bundled instance also syncs its light/dark theme to match the user's chosen SO-CRATES theme unless a CyberChef theme has been set manually.

— Explains size and offline mechanism but no config surface to act on v4.3.0

04 **Send to CyberChef for Suricata-extracted files** NEW 50

Adds a Send to CyberChef button in the File Info panel for files Suricata extracted from traffic, loading the file's bytes directly into the bundled CyberChef instance for offline triage.

Inspect a file Suricata carved from traffic — PE, archive, document — directly in CyberChef for offline triage.

📍 Navigate to the alert or flow's File Info panel and click the Send to CyberChef button next to the extracted file to load its bytes into the bundled CyberChef instance.

— Names the panel and button but describes little beyond the single action v4.3.0

v4.3.0 NOTES ↗ LATEST STABLE

SO-CRATES v4.3.0 ships a bundled, air-gapped CyberChef with full-payload pivoting, bidirectional stream support, and Suricata-extracted file analysis.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v4.3.0 https://github.com/dougburks/so-crates.git
# already have the repo? check out this version:
$ git checkout v4.3.0
```

└──▷ HOW TO FIND IT

Analyse a full bidirectional TCP payload — shellcode, compressed data, XOR-keyed content — in an air-gapped CyberChef without it appearing in browser history.

📍 In the stream detail view, open the Payload panel, then click Send to CyberChef and choose Both, Source, or Dest to load the exact bytes into the bundled CyberChef instance.

Decode a suspicious base64 blob or header value spotted in an ASCII transcript without leaving the SO-CRATES UI.

📍 Select the text (e.g. a base64 blob) in the ASCII transcript, then choose CyberChef from the pivot menu that appears — CyberChef opens with the selected text and, for small values, applies the Magic operation automatically.

Inspect a file Suricata carved from traffic — PE, archive, document — directly in CyberChef for offline triage.

📍 Navigate to the alert or flow's File Info panel and click the `Send to CyberChef` button next to the extracted file to load its bytes into the bundled CyberChef instance.

Bundles CyberChef (~48 MB added to container image) so the pivot menu's CyberChef lookup works with no internet access.

Adds `Send to CyberChef` from a stream's Payload panel, sending exact bytes in Both directions, or just the Source or Dest side — XOR keys, compressed data, and shellcode arrive intact.

Adds a `Send to CyberChef` button in File Info for files Suricata extracted from traffic.

Enables pivot-menu access to CyberChef (and other lookups) by selecting any text in an ASCII transcript — base64 blobs, headers, or partial values.

Applies CyberChef's Magic operation automatically to small payloads, now including values sent from the pivot menu, to suggest decodings.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ

Add AI Firewall metrics to Historical Stats and Real-Time Analytics APIs (2026-09-23)

BLOG ↗ RANK

Fastly's edge cloud platform delivers faster, safer, and more scalable sites and apps to customers. Elevate your edge CDN, video delivery, security, and more.

> Fastly expanded its analytics APIs with new AI Firewall and Content Guard/PAT metrics, and overhauled Next-Gen WAF event logging to capture changes from every management surface while deprecating the old event-type taxonomy.

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

01 Next-Gen WAF changes recorded across all surfaces in event log

95

BREAKING

The Event logs API now records Next-Gen WAF configuration changes made through any surface — the Next-Gen WAF API, the Fastly control panel, the Next-Gen WAF control panel, or Terraform — not just changes made through Fastly. The list events API gains `filter[event_type]` support for all `ngwaf.*` event types, covering workspaces, rules, lists, custom signals, redactions, virtual patches, thresholds, workspace alerts, header links, rate limited sources, custom dashboards, request reports, and agent keys. The legacy `security.workspace.create`, `security.workspace.update`, `security.workspace.delete`, `security.redaction.create`, `security.redaction.update`, `security.redaction.delete`, `security.rule.create`, `security.rule.update`, `security.rule.delete`, `security.virtual_patch.create`, `security.virtual_patch.update`, `security.virtual_patch.delete`, and `security.event.update` event types are deprecated and no longer recorded — use the `ngwaf.*` equivalents (e.g. `ngwaf.rule.create` in place of `security.rule.create`); existing log entries keep their original event type.

Audit all Next-Gen WAF rule changes across every management surface — API, control panels, and Terraform — by filtering the event log for ngwaf.rule events.

```
$ curl -s -X GET 'https://api.fastly.com/events?
filter[event_type]=ngwaf.rule.create' -H 'Fastly-Key:
<your_api_token>'
```

Review all Next-Gen WAF workspace deletions to detect unauthorised removal of workspaces regardless of which panel or tool made the change.

```
$ curl -s -X GET 'https://api.fastly.com/events?
filter[event_type]=ngwaf.workspace.delete' -H 'Fastly-Key:
<your_api_token>'
```

— Exact endpoint, filter syntax, full deprecated-type list, runnable curl examples.

Next-Gen WAF configuration changes now …

THINNER COVERAGE BELOW

02 **AI Firewall metrics in analytics APIs** NEW 55

The Historical Stats API and Real-Time Analytics API gain two new metrics, `arc_ai_firewall_requests` and `arc_ai_firewall_blocked`, tracking AI Firewall activity within AI Runtime Control.

— Names exact metrics but gives no query example or endpoint.

Add AI Firewall metrics to Historical S…

03 **Content Guard and PAT metrics in analytics APIs** NEW 40

Adds ten new metrics to the Historical Stats API and Real-Time Analytics API covering Content Guard challenge and token creation, validation outcomes, and Private Access Token (PAT) challenges issued during rule enforcement.

— Describes scope but omits the actual metric names.

Add Content Guard challenge metrics to …

└──▷ **BREAKING ON UPGRADE**

! The event types `security.workspace.create`, `security.workspace.update`, `security.workspace.delete`, `security.redaction.create`, `security.redaction.update`, `security.redaction.delete`, `security.rule.create`, `security.rule.update`, `security.rule.delete`, `security.virtual_patch.create`, `security.virtual_patch.update`, `security.virtual_patch.delete`, and `security.event.update` are deprecated and no longer recorded — use the `ngwaf.*` equivalents (e.g. `ngwaf.rule.create` in place of `security.rule.create`); existing log entries keep their original event type.

> Add AI Firewall metrics to Historical Stats and Real-Time Analytics APIs (2026-09-23)
>
> BLOG ↗ STABLE
>
> Fastly Historical Stats and Real-Time Analytics APIs gain two new AI Firewall metrics for AI Runtime Control.

Next-Gen WAF configuration changes now recorded in the event log (2026-09-08)

BLOG ↗ STABLE

Event logs API now captures Next-Gen WAF config changes from all sources — API, Fastly panel, NG WAF panel, and Terraform.

└──▷ **TRY IT**

Audit all Next-Gen WAF rule changes across every management surface — API, control panels, and Terraform — by filtering the event log for ngwaf.rule events.

```
$ curl -s -X GET 'https://api.fastly.com/events?
filter[event_type]=ngwaf.rule.create' -H 'Fastly-Key:
<your_api_token>'
```

Review all Next-Gen WAF workspace deletions to detect unauthorised removal of workspaces regardless of which panel or tool made the change.

```
$ curl -s -X GET 'https://api.fastly.com/events?
filter[event_type]=ngwaf.workspace.delete' -H 'Fastly-Key:
<your_api_token>'
```

Adds `filter[event_type]` support on the list events API for all `ngwaf.*` event types, covering workspaces, rules, lists, custom signals, redactions, virtual patches, thresholds, workspace alerts, header links, rate limited sources, custom dashboards, request reports, and agent keys.

Extends the Event logs API to record Next-Gen WAF configuration changes made through any surface — the Next-Gen WAF API, the Fastly control panel, the Next-Gen WAF control panel, or Terraform — not just changes made through Fastly.

└──▷ **BREAKING ON UPGRADE**

! The event types `security.workspace.create`, `security.workspace.update`, `security.workspace.delete`, `security.redaction.create`, `security.redaction.update`, `security.redaction.delete`, `security.rule.create`, `security.rule.update`, `security.rule.delete`, `security.virtual_patch.create`, `security.virtual_patch.update`, `security.virtual_patch.delete`, and `security.event.update` are deprecated and no longer recorded — use the `ngwaf.*` equivalents (e.g. `ngwaf.rule.create` in place of `security.rule.create`); existing log entries keep their original event type.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Sep 29, 2026 | DOCS ↗ | BLOG ↗ | RANK

Rapid7 InsightCloudSec is a cloud security platform that provides visibility, compliance monitoring, and vulnerability management across cloud infrastructures.

> InsightCloudSec's September 2026 release adds CSV export for container vulnerability findings, two new CIS database compliance packs for AWS and Azure, two new RDS misconfiguration detections, and small improvements to Layered Context and OCI harvesting.

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

01 CIS database compliance packs for AWS and Azure NEW

75

Adds `CIS AWS Database Services Benchmark v1.0.0` compliance pack covering 22 controls across Aurora, RDS, DynamoDB, ElastiCache, Redshift, DocumentDB, Neptune, and Timestream. Adds `CIS Microsoft Azure Database Services Benchmark v2.0.0` compliance pack covering 47 controls across Azure Cache for Redis, Azure Cosmos DB, Azure Data Factory, Azure Database for MySQL, Azure Database for PostgreSQL, and Azure SQL Database — available for Azure, Azure Gov, and Azure China.

— Names both packs, control counts, and covered services 2026-09-29-023b1562

02 Container vulnerability CVE CSV export NEW

65

Exports container vulnerability findings from the Layered Context list view as a CSV file, with one row per container, CVE, and affected package.

Pull a per-container, per-CVE, per-package CSV of all container vulnerability findings for offline triage or ticket import.

> 📍 In InsightCloudSec, open Layered Context, filter to the desired container resources, then select Export › CSV to download one row per container, CVE, and affected package.

— Names export format and row structure, gives UI steps 2026-09-29-023b1562

03 Database transit encryption and backup detections NEW

65

Adds Database Instance Not Enforcing Transit Encryption (AWS) detection identifying Amazon RDS instances that do not enforce SSL/TLS encryption in transit, excluding

DocumentDB and Neptune, and Database Instance Without Automatic Backups (AWS) detection identifying Amazon RDS instances with backup retention period of 0, excluding DocumentDB and Neptune.

— Names exact detection conditions and exclusions 2026-09-29-023b1562

THINNER COVERAGE BELOW

04 **Reduced OCI Identity Domains API calls** IMPROVED 45

Reduces OCI Identity Domains API calls during cloud user harvesting for Oracle Cloud Infrastructure, lowering the risk of rate-limit errors on large tenancies.

— Explains benefit but no user-facing action 2026-09-29-023b1562

05 **Resource lifecycle state in Layered Context** IMPROVED 40

Shows each resource's lifecycle state as reported by the cloud provider, along with last harvest time, in Layered Context.

— Brief description, no mechanism beyond display 2026-09-29-023b1562

changelog · Sep 29, 2026 DOCS ↗ BLOG ↗ LATEST STABLE

InsightCloudSec 26.9.29 adds CIS database compliance packs, container CVE CSV export, and richer Layered Context resource state.

└──▷ **HOW TO FIND IT**

Pull a per-container, per-CVE, per-package CSV of all container vulnerability findings for offline triage or ticket import.

📍 In InsightCloudSec, open Layered Context, filter to the desired container resources, then select Export › CSV to download one row per container, CVE, and affected package.

Exports container vulnerability findings from the Layered Context list view as a CSV file, with one row per container, CVE, and affected package.

Adds `CIS AWS Database Services Benchmark v1.0.0` compliance pack covering 22 controls across Aurora, RDS, DynamoDB, ElastiCache, Redshift, DocumentDB, Neptune, and Timestream.

Adds `CIS Microsoft Azure Database Services Benchmark v2.0.0` compliance pack covering 47 controls across Azure Cache for Redis, Azure Cosmos DB, Azure Data Factory, Azure Database for MySQL, Azure Database for PostgreSQL, and Azure SQL Database — available for Azure, Azure Gov, and Azure China.

Adds Database Instance Not Enforcing Transit Encryption (AWS) detection identifying Amazon RDS instances that do not enforce SSL/TLS encryption in transit,

> excluding DocumentDB and Neptune.
>
> Adds Database Instance Without Automatic Backups (AWS) detection identifying Amazon RDS instances with backup retention period of 0, excluding DocumentDB and Neptune.
>
> +2 more

WAS THIS USEFUL?  

◆ Container & Image Security

kubescape

Feedback ⓘ v4.0.15 NOTES ↗ RANK

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

> Kubescape v4.0.15 adds multi-cluster fleet scanning with a combined report and compliance rollup, expands image scanning to Quay and Azure ACR credential chains, adds scan-coverage metrics to report outputs, and speeds up large-cluster processing with a new partition and spill store engine.

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

01 Multi-cluster fleet scanning NEW

90

The `--fleet-report` flag on `kubescape scan`, combined with `--kube-contexts` to name multiple clusters, writes one unified report across all of them. The fleet report includes a compliance rollup, a fleet summary printed after the scan, and reporting of which controls clusters disagree on.

Scan multiple Kubernetes clusters in one pass and write a unified fleet report — useful in multi-tenant or multi-region environments.

```
$ kubescape scan --kube-contexts prod-us,prod-eu,staging --fleet-report fleet-results.json
```

— Exact flags plus a runnable multi-cluster command v4.0.15

02 Show-evidence fix-path line resolution IMPROVED

60

The `--show-evidence` output now resolves fix-path line numbers for delete and review paths.

— Names the flag and behaviour, no example given v4.0.15

03 Config command key filtering and deletion NEW

60

`kubescape config view` supports filtering configurations by key, and `kubescape config` now supports deleting specific configuration keys.

— Two named subcommand capabilities, no usage shown v4.0.15

04 Scan coverage metrics and degraded-coverage reporting NEW

60

Adds scan coverage metrics to Markdown, HTML, and CSV report outputs, and reports degraded scan coverage and skipped controls in printer output.

— Names three report formats but no command shown v4.0.15

05 **Large-cluster performance improvements** IMPROVED 60

Introduces a partition and spill store engine for handling large-cluster resources, and uses mmap and fastjson for manifest parsing to improve performance on large manifest sets.

— Names the mechanisms but no benchmark or command v4.0.15

THINNER COVERAGE BELOW

06 **VAP Agent Runtime policy set** NEW 55

`kubescape vap` gains a complete Agent Runtime policy set.

— Names exact subcommand, scope of policy set unclear v4.0.15

07 **Key-signed image support without transparency log** NEW 55

Control `C-0236` now supports key-signed images that lack a transparency log entry.

— Names exact control ID, no further mechanism detail v4.0.15

08 **Output-dir support for file-based reports** IMPROVED 45

The `--output-dir` flag is now honoured for file-based reports.

— Bare one-line addition naming an exact flag v4.0.15

09 **Quay and Azure ACR image scanning support** NEW 45

Image scanning now supports Quay vulnerability scanning, and full Azure credential chain support for ACR image pulls.

— Names two registries, no mechanism or flags given v4.0.15

10 **Namespace filter hot-reload** IMPROVED 40

Namespace filters in the HTTP handler now hot-reload without requiring a restart.

— Explains behaviour but no config or endpoint named v4.0.15

11 **OPA whole-cluster execution policies** NEW 40

Adds execution policies for whole-cluster controls in the OPA processor.

— Brief mention of processor change, scope unclear v4.0.15

v4.0.15 NOTES ↗ LATEST STABLE

Kubescape v4.0.15 adds multi-cluster fleet scanning, Quay image scanning, namespace-filter hot-reload, and scan-coverage metrics in reports.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v4.0.15
https://github.com/kubescape/kubescape.git
# already have the repo? check out this version:
$ git checkout v4.0.15
```

└──▷ **TRY IT**

Scan multiple Kubernetes clusters in one pass and write a unified fleet report — useful in multi-tenant or multi-region environments.

```
$ kubescape scan --kube-contexts prod-us,prod-eu,staging --fleet-
report fleet-results.json
```

Adds `--fleet-report` flag to `kubescape scan` to write one combined report across multiple clusters specified with `--kube-contexts`.

Adds `--show-evidence` output that now resolves fix-path line numbers for delete and review paths.

Adds `kubescape config view` support for filtering configurations by key.

Adds support for deleting specific configuration keys via `kubescape config`.

Adds `--output-dir` support honoured for file-based reports.

+13 more

WAS THIS USEFUL?  

Chainguard

Feedback ⓘ Week of 2026-09-28 NOTES ↗ RANK

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

> Chainguard shipped a FIPS 140-3 validated OpenSSL provider with post-quantum key exchange, opened an invite-only beta for Chainguard Libraries covering Go and .NET, added 19 new container images, and expanded the Agent Skills CLI with new publishing, install and lifecycle flags.

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

01 Agent Skills CLI flags and lifecycle handling IMPROVED

95

The `chainctl skills` CLI adds `--dry-run` to the publish command to build and validate an artifact without publishing it; `--tag` (repeatable, e.g. `--tag v1.0.0 --tag latest`) to publish under one or more explicit tags, warning when omitted that no version was pinned; `--copy` on the install command to give each agent its own copy of a skill instead of the default shared symlink under `.agents/skills/`; `--yes` to non-interactively accept the Agent Skills Terms of Service and Data Privacy Agreement in CI environments; and `--parent` on `entitlements` subcommands plus `--group` on other `skills` subcommands to target a specific organization. Deleting a skill's last remaining version now also removes the empty skill entry to prevent stale listing output, and skill uninstall accepts a tag-less reference (e.g. `"$HARDENED_REPO"`) only when no versions remain, otherwise an explicit tag is required. Requires `chainctl` v0.2.364 or later.

Validate a skill artifact locally before publishing — useful in PR CI to catch build errors without creating a registry entry.

```
$ chainctl skills publish --tag v1.0.0 --dry-run
```

Publish the same skill artifact under both a pinned version tag and `latest` so consumers can reference either.

```
$ chainctl skills publish --tag v1.0.0 --tag latest
```

— Names every flag with runnable example commands product docs

02 FIPS 140-3 validated OpenSSL provider with post-quantum support

NEW

85

Chainguard FIPS Provider for OpenSSL v3.6 is FIPS 140-3 validated (Certificate #5523), adding hybrid post-quantum TLS key exchange via ML-KEM paired with classical algorithms, plus standalone ML-KEM, ML-DSA, and SLH-DSA support alongside AES, SHA-2, SHA-3, ECDSA, and RSA. Opt-in is available for Chainguard FIPS container images starting October 1, 2026; note that SHA-1 is not an approved algorithm in the v3.6 FIPS module, so workloads relying on SHA-1 will break if they opt in on or after that date.

— Names cert number and algorithms but no opt-in command shown Week of 2026-09-28

03 **19 new container images added** NEW

60

Adds 19 new container images this week, including `chainguard-server-hypervisor-gcp`, `kured` (+fips), `conftest` (+fips), `foundry`, `keepalived` (+fips), `mcp-grafana-fips`, `nix`, `ray`, `vault-secrets-webhook` (+fips), `argocd-agent`, `kates` (+fips), `kubernetes-pause-fips`, and `victoriatraces` (+fips).

— Full image list named but no usage detail beyond naming Week of 2026-09-28

THINNER COVERAGE BELOW

04 **Chainguard Libraries beta for Go and .NET** NEW

45

Chainguard Libraries launches for Go and .NET ecosystems in an invite-only beta, letting accepted organizations pull Go and .NET packages from Chainguard Libraries into existing development workflows.

— Invite-only beta described with no usage steps Week of 2026-09-28

└──▷ **BREAKING ON UPGRADE**

! SHA-1 is not an approved algorithm in the OpenSSL v3.6 FIPS module — workloads relying on SHA-1 will break if they opt in to the v3.6 module on or after October 1, 2026.

Week of 2026-09-28 NOTES ↗ STABLE

Chainguard gains FIPS 140-3 validation with post-quantum algorithms, Libraries beta for Go and .NET, and 19 new container images.

Adds FIPS 140-3 validated Chainguard FIPS Provider for OpenSSL v3.6 (Certificate #5523), including hybrid post-quantum TLS key exchange via ML-KEM paired with classical algorithms, plus standalone ML-KEM, ML-DSA, and SLH-DSA support alongside AES, SHA-2, SHA-3, ECDSA, and RSA; opt-in available for Chainguard FIPS container images starting October 1, 2026.

Launches Chainguard Libraries for Go and .NET ecosystems in an invite-only beta, letting accepted organizations pull Go and .NET packages from Chainguard Libraries into existing development workflows.

Adds 19 new container images this week, including `chainguard-server-hypervisor-gcp`, `kured` (+fips), `conftest` (+fips), `foundry`, `keepalived` (+fips), `mcp-grafana-fips`, `nix`, `ray`, `vault-secrets-webhook` (+fips), `argocd-agent`, `kates` (+fips), `kubernetes-pause-fips`, and `victoriatraces` (+fips).

└──▷ **BREAKING ON UPGRADE**

! SHA-1 is not an approved algorithm in the OpenSSL v3.6 FIPS module — workloads relying on SHA-1 will break if they opt in to the v3.6 module on or after October 1, 2026.

WAS THIS USEFUL? 👍 👎

Anchore

Feedback ⓘ

FedRAMP & CMMC in 2026: what actually changed for your Authority to Operate (ATO) (2026-

BLOG ↗ RANK

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

> Anchore's policy bundle v2026.1 adds a CISA KEV-mapped STOP gate with per-control failure detail, folds KEV and EPSS scoring into risk-based prioritization, adds POA&M-style time-bound allowlisting, and extends policy enforcement from build time into live Kubernetes clusters.

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

01 STOP gate for CISA KEV vulnerabilities with per-control detail

NEW 87

Policy bundle v2026.1 adds a dedicated STOP gate for CISA KEV-listed vulnerabilities mapped to NIST 800-53 SI-2. Running `anchorectl image check <image> --detail -f` surfaces per-control failure detail, returning the specific control and finding that failed rather than a pass/fail count, so assessors can trace results back to a named requirement.

Fail a CI/CD build immediately on any CISA KEV-listed vulnerability and get the specific NIST control and finding — not just a pass/fail — so the result is directly traceable for an ATO assessor.

```
$ anchorectl image check <image> --detail -f
```

— Names the gate, control mapping, and exact command with flags

FedRAMP & CMMC in 2026: what actually c…

THINNER COVERAGE BELOW

02 Runtime policy enforcement in Kubernetes via admission controller

NEW 52

Extends policy enforcement from build time into live Kubernetes clusters via a runtime inventory and admission controller that blocks non-compliant workloads before scheduling.

— States mechanism and scope but no setup steps given

FedRAMP & CMMC in 2026: what actually c…

03 **POA&M enforcement with time-bound allowlists** NEW 49

Supports POA&M enforcement through policies, gates, and time-bound allowlists, letting a non-compliant finding be allowlisted with an explicit remediation date rather than silently ignored or blocking all builds.

— Describes mechanism but no concrete config key or command

FedRAMP & CMMC in 2026: what actually c…

04 **KEV and EPSS scoring in policy engine** NEW 46

Incorporates CISA's Known Exploited Vulnerabilities (KEV) list and EPSS exploit prediction scores alongside CVSS in the policy engine, enabling risk-based vulnerability prioritization.

— Names data sources but no config or command shown

FedRAMP & CMMC in 2026: what actually c…

FedRAMP & CMMC in 2026: what actually changed for your Authority to Operate (ATO) (202

BLOG ↗ LATEST STABLE

Anchore Enterprise policy bundle v2026.1 adds a dedicated STOP gate for CISA KEV-listed vulnerabilities mapped to NIST 800-53 SI-2.

└──▷ **TRY IT**

Fail a CI/CD build immediately on any CISA KEV-listed vulnerability and get the specific NIST control and finding — not just a pass/fail — so the result is directly traceable for an ATO assessor.

```
$ anchorectl image check <image> --detail -f
```

Surfaces per-control failure detail via `anchorectl image check <image> --detail -f`, returning the specific control and finding that failed rather than a pass/fail count, enabling assessors to trace results to a named requirement.

Incorporates CISA's Known Exploited Vulnerabilities (KEV) list and EPSS exploit prediction scores alongside CVSS in the policy engine for risk-based vulnerability prioritization.

Supports POA&M enforcement through policies, gates, and time-bound allowlists, letting a non-compliant finding be allowlisted with an explicit remediation date rather than silently ignored or blocking all builds.

Extends policy enforcement from build time into live Kubernetes clusters via a runtime inventory and admission controller that blocks non-compliant workloads before scheduling.

WAS THIS USEFUL?  

Snyk Container

Feedback ⓘ changelog · Sep 16, 2026 BLOG ↗

RANK

Snyk Container scans Docker images and container registries for vulnerabilities, misconfigurations, and malware to secure containerized applications.

> Snyk Container now detects and reports vulnerabilities in the Go standard library itself, automatically surfacing them as a new dependency in scan results.

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

01 **Go standard library vulnerability reporting** NEW 78

Snyk Container scans now identify the Go standard library version a binary was compiled with and surface known CVEs against it, shown as a `stdlib@<version>` dependency (e.g. `stdlib@1.25.10`) in the dependency graph. It works across standard, stripped, and CGo builds, is enabled automatically for all organizations with no configuration required, and CLI-based container scans require Snyk CLI v1.1303.2 or later. As a result, vulnerability counts will increase for container image projects built from Go binaries on or after October 7, 2026, except those built with Go 1.25.13+ or Go 1.26.6+.

— Names format, versions, CLI requirement, and rollout impact date. 2026-09-16-523bbbee

└──▷ **BREAKING ON UPGRADE**

! Vulnerability counts will increase for all container image projects built from Go binaries on or after October 7, 2026: each affected project gains new findings determined by the Go version the binary was compiled with. Images built with Go 1.25.13+ or Go 1.26.6+ report no new findings from this change.

changelog · Sep 16, 2026 BLOG ↗ LATEST STABLE

Snyk Container now reports Go standard library vulnerabilities in container images, shown as a `stdlib@<version>` dependency.

Adds Go standard library vulnerability reporting for container images: scans identify the stdlib version the binary was compiled with and surface known CVEs against it, shown as `stdlib@<version>` (e.g. `stdlib@1.25.10`) in the dependency graph — enabled automatically for all organizations with no configuration required.

Supports Go standard library detection across standard, stripped, and CGo builds.

Requires Snyk CLI v1.1303.2 or later for standard library reporting in CLI-based container scans.

└──▷ **BREAKING ON UPGRADE**

! Vulnerability counts will increase for all container image projects built from Go binaries on or after October 7, 2026: each affected project gains new findings

determined by the Go version the binary was compiled with. Images built with Go 1.25.13+ or Go 1.26.6+ report no new findings from this change.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.44.0 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 5.44.0 adds targeted partial re-scans of specific resources across the API, MCP tools and UI, plus S3-compatible output storage and configurable AWS retry behavior.

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

01 Partial resource re-checks via API, MCP and UI NEW

95

Adds `POST /api/v1/scans` support for `resource_uids` (up to 10 resources) to launch partial scans, with `is_partial` exposed on scan objects and a `filter[is_partial]` query parameter on scan list endpoints. The same capability is exposed through the `prowler_trigger_scan` MCP tool's `resource_uids` argument, with `is_partial` shown on `prowler_list_scans` and `prowler_get_scan` and an `is_partial` filter on `prowler_list_scans`. A new Re-check resource action in the Findings page actions menu and resource detail drawer launches a partial scan that re-runs only the checks that last reported on that resource.

Immediately re-check up to 10 specific resources after a remediation, without waiting for the next full scan.

```
$ curl -X POST https://<your-prowler-api>/api/v1/scans \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"resource_uids": ["arn:aws:s3:::my-bucket",
"arn:aws:ec2:us-east-1:123456789012:instance/i-0abc123"]}'
```

— Names endpoint, fields, filter, MCP tool and UI action verbatim. 5.44.0

02 S3-compatible storage for scan output NEW

80

Adds `DJANGO_OUTPUT_S3_AWS_ENDPOINT_URL` environment variable to direct scan output uploads and downloads at S3-compatible object storage such as MinIO, without polluting process-wide AWS environment variables.

Point a self-hosted Prowler deployment at MinIO (or another S3-compatible store) for scan report storage without affecting AWS provider role assumptions.

```ini
DJANGO_OUTPUT_S3_AWS_ENDPOINT_URL=http://minio.internal:9000
```

— Exact env var and config example given. 5.44.0

03 **Configurable Boto3 retry attempts for AWS provider** NEW

75

Adds `PROWLER_AWS_BOTO3_RETRIES_MAX_ATTEMPTS` environment variable to control Boto3 retry attempts for the AWS provider; set to `0` to disable retries entirely.

Cap Boto3 retries for the AWS provider in environments where unreachable partition endpoints cause long delays during connection tests.

```ini
PROWLER_AWS_BOTO3_RETRIES_MAX_ATTEMPTS=3
```

— Exact env var, default behavior and disable value named. 5.44.0

THINNER COVERAGE BELOW

04 **Simplified AWS account connection wizard** IMPROVED

55

Connects AWS accounts in a single wizard step: the account ID is read from the role ARN (or typed for static access keys), the role is assumed, and credentials are stored and tested in one submit.

— Describes flow mechanism but no exact UI path or endpoint. 5.44.0

05 **Unified latest-scan resolution across endpoints** IMPROVED

35

Unifies latest-scan resolution across all endpoints so findings, compliance, and mute rules always read from the same scan for a given provider.

— No named endpoints or config; thin prose description. 5.44.0

5.44.0 NOTES ↗ CODE ↗ LATEST STABLE

Prowler 5.44.0 adds partial resource re-checks, S3-compatible storage via `DJANGO_OUTPUT_S3_AWS_ENDPOINT_URL`, and `PROWLER_AWS_BOTO3_RETRIES_MAX_ATTEMPTS` for retry control.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 5.44.0 https://github.com/prowler-cloud/prowler.git
# already have the repo? check out this version:
$ git checkout 5.44.0
```

└──▷ **TRY IT**

Immediately re-check up to 10 specific resources after a remediation, without waiting for the next full scan.

```
$ curl -X POST https://<your-prowler-api>/api/v1/scans \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"resource_uids": ["arn:aws:s3:::my-bucket", "arn:aws:ec2:us-east-1:123456789012:instance/i-0abc123"]}'
```

Point a self-hosted Prowler deployment at MinIO (or another S3-compatible store) for scan report storage without affecting AWS provider role assumptions.

.env

```
DJANGO_OUTPUT_S3_AWS_ENDPOINT_URL=http://minio.internal:9000
```

Cap Boto3 retries for the AWS provider in environments where unreachable partition endpoints cause long delays during connection tests.

.env

```
PROWLER_AWS_BOTO3_RETRIES_MAX_ATTEMPTS=3
```

Adds `POST /api/v1/scans` support for `resource_uids` (up to 10 resources) to launch partial scans, with `is_partial` exposed on scan objects and a `filter[is_partial]` query parameter on scan list endpoints.

Adds `DJANGO_OUTPUT_S3_AWS_ENDPOINT_URL` environment variable to direct scan output uploads and downloads at S3-compatible object storage such as MinIO, without polluting process-wide AWS environment variables.

Adds `PROWLER_AWS_BOTO3_RETRIES_MAX_ATTEMPTS` environment variable to control Boto3 retry attempts for the AWS provider; set to `0` to disable retries entirely.

Adds `prowler_trigger_scan` MCP tool support for a `resource_uids` argument to launch partial scans, and exposes `is_partial` on `prowler_list_scans` and `prowler_get_scan` with an `is_partial` filter on `prowler_list_scans`.

> Adds a `Re-check resource` action in the Findings page actions menu and resource detail drawer, launching a partial scan that re-runs only the checks that last reported on that resource.
>
> +2 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Sumo Logic

Feedback ⓘ Amazon Security Lake (Apps) (2026-09-29)

BLOG ↗ RANK

Sumo Logic is a cloud-based platform for log management, monitoring, and security analytics of applications and infrastructure.

> Sumo Logic added a new app for Amazon Security Lake, giving unified AWS security visibility through OCSF-normalized events and pre-built dashboards.

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

01 **Amazon Security Lake app** NEW 45

A new Sumo Logic app for Amazon Security Lake provides unified visibility into AWS security activity across accounts and regions using OCSF-normalized events, with pre-built dashboards covering network, identity, application, and host activity to detect threats, investigate anomalies, and assess AWS security posture.

— Describes scope and dashboards but no setup steps or config details.

Amazon Security Lake (Apps) (2026-09-29)

> Amazon Security Lake (Apps) (2026-09-29) BLOG ↗ LATEST STABLE
>
> New Sumo Logic app for Amazon Security Lake delivers unified AWS security visibility via OCSF-normalized events and pre-built dashboards.
>
> Adds a new Sumo Logic app for Amazon Security Lake that provides unified visibility into AWS security activity across accounts and regions using OCSF-normalized events.
>
> Includes pre-built dashboards covering network, identity, application, and host activity to detect threats, investigate anomalies, and assess AWS security posture.

WAS THIS USEFUL? 👍 👎

Databricks Feedback ⓘ

1Password Event Logs connector (Beta) (2026-09-28) BLOG ↗ STABLE RANK

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

Lakeflow Connect gains a 1Password Event Logs connector (Beta); workspace admins control access via the Previews page.

NEW INTEGRATION Adds a 1Password Event Logs connector in Lakeflow Connect (Beta), ingesting audit events, item usages, and sign-in attempts from the 1Password Events API into Databricks using an Events Reporting bearer token.

Workspace admins can enable or disable the 1Password Event Logs connector from the Previews page.

NEW INTEGRATION Workday Activity Logging connector

▾ 4 more releases in this issue · 2026-09-24 → 2026-09-28

> Microsoft 365 Unified Audit Logs connector is available in Beta (2026-09-28) BLOG ↗ STABLE
>
> Databricks Lakeflow Connect adds a managed Microsoft 365 Unified Audit Logs connector (Beta) using OAuth machine-to-machine auth with incremental ingestion.

> Workday Activity Logging connector (Beta) (2026-09-28) BLOG ↗ LATEST STABLE
>
> Databricks Lakeflow Connect introduces a Beta Workday Activity Logging connector that incrementally ingests Workday user activity log events using OAuth refresh tokens.

> Usage tracking adds session metadata and token details (Public Preview) (2026-09-27) BLOG ↗ STABLE
>
> system.ai_gateway.usage gains coding-agent session metadata, invocation source, prompt-cache token fields, and search counts.
>
> Adds token-detail fields to `system.ai_gateway.usage` capturing tokens written to the 5-minute and 1-hour prompt caches, plus file and web search counts.

> ai_search can search Lakebase synced tables (Beta) (2026-09-24) BLOG ↗ STABLE

> ai_search now supports Lakebase synced tables via a new `lakebase_search` knowledge source, combinable with `vector_search` in one call (Beta).

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Eclypsium

Feedback ⓘ

Safe AI Agents Start With Trusted AI Infrastructure (2026-09-28) BLOG ↗ RANK

Eclypsium's infrastructure assurance platform protects mission critical hardware and detects compromised devices in enterprises and government agencies.

> Eclypsium launched AI Infrastructure Assurance, joining NVIDIA's Open Agent Safety Platform to extend hardware and firmware integrity monitoring to the GPUs, DPUs, BMCs and other components underpinning AI agent workloads.

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

01 **NVIDIA Open Agent Safety Platform integration** NEW

60

Integrates with NVIDIA's Open Agent Safety Platform to extend Eclypsium's firmware integrity verification and vulnerability identification to the infrastructure underpinning NVIDIA OpenShell and NVIDIA Sentry enforcement boundaries. Continuously monitors AI infrastructure components for integrity drift, unauthorized firmware modifications, and persistent compromise throughout the device lifecycle, and provides component-level evidence — including known-good baseline comparisons and supply-chain risk — to inform decisions about when agent workloads are permitted to run.

— Explains mechanism and lifecycle but no concrete config or command

Safe AI Agents Start With Trusted AI In…

THINNER COVERAGE BELOW

02 **AI Infrastructure Assurance coverage for AI hardware** NEW

55

Adds AI Infrastructure Assurance coverage that inventories hardware, firmware, and OS code across AI-specific components — GPUs, DPUs (including NVIDIA BlueField-4), TPUs, BMCs, NICs, HBM, SSDs, TPMs, and accelerators.

— Names many components but no command or UI path

Safe AI Agents Start With Trusted AI In…

Safe AI Agents Start With Trusted AI Infrastructure (2026-09-28) BLOG ↗

LATEST STABLE

Eclypsium joins NVIDIA's Open Agent Safety Platform to extend hardware/firmware assurance to AI agent infrastructure including GPUs, DPUs, and BMCs.

> Adds AI Infrastructure Assurance coverage for AI-specific components — GPUs, DPUs (including NVIDIA BlueField-4), TPUs, BMCs, NICs, HBM, SSDs, TPMs, and accelerators — inventorying hardware, firmware, and OS code across all of them.
>
> Integrates with NVIDIA Open Agent Safety Platform to extend Eclypsium's firmware integrity verification and vulnerability identification to the infrastructure underpinning NVIDIA OpenShell and NVIDIA Sentry enforcement boundaries.
>
> Continuously monitors AI infrastructure components for integrity drift, unauthorized firmware modifications, and persistent compromise throughout the device lifecycle — not just at deployment time.
>
> Provides component-level evidence of device and firmware posture — including known-good baseline comparisons and supply-chain risk — to inform decisions about when agent workloads are permitted to run.

WAS THIS USEFUL?  

Trend Vision One

Feedback ⓘ

Scan streaming, cookie-authenticated, and plain-text custom AI endpoints in AI Scanner (

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 expanded AI Scanner to handle streaming, cookie-authenticated, and plain-text custom AI endpoints, rolled out ZTSA module improvements for Android/ChromeOS, and moved Agentless Vulnerability & Threat Detection on Azure to the Flex Consumption Plan while adding China region coverage and tighter transport security.

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

01 **Streaming, cookie auth, and plain-text support for custom AI endpoints**

NEW 95

AI Scanner adds a `stream` option for custom targets so it can scan endpoints returning incremental responses (SSE or streamed text chunks); in JSON streaming mode it applies the response template to each frame as a delta path (e.g. `choices[0].delta.content`) and concatenates content in arrival order, while text streaming mode reads and joins raw bytes until the stream ends. It also adds `cookie` as a supported `auth_type` value, reading credentials from `api_key_env` and sending them in the `cookie` header, and adds `text` as a supported `response_format` value that treats the whole response body as raw text (requiring the `{{response}}` marker in the response template); plain-text handling works with both buffered and streaming endpoints and can be combined with cookie auth or streaming.

Scan a cookie-authenticated, streaming chat endpoint that returns plain-text chunks — combining all three new capabilities in one custom target config.

yaml

```yaml
auth_type: cookie
api_key_env: SESSION_COOKIE
stream: true
response_format: text
response_template: '{{response}}'
```

Scan a streaming JSON endpoint (e.g. an OpenAI-compatible API) by enabling the `stream` option and pointing the delta path at the incremental content field.

```yaml
stream: true
response_format: json
response_template: 'choices[0].delta.content'
```

— Names exact config keys, values and runnable example configs

Scan streaming, cookie-authenticated, a…

02 **China region support for Agentless V&TD** NEW 75

Adds China East 2, China North 2, and China North 3 region support for Agentless Vulnerability & Threat Detection and VNet Flow Logs, including vulnerability and malware scanning for Azure Virtual Machines, Disks, Azure Container Registry images, and Azure Function Apps, plus VNet Flow Logs collection and forwarding to TrendAI Vision One.

— Names exact regions and covered Azure resource types

Agentless Vulnerability & Threat Detect…

03 **Azure Flex Consumption Plan migration for Agentless V&TD**

IMPROVED 65

Agentless Vulnerability & Threat Detection and VNet Flow Logs are migrated from the Y1 Consumption Plan to the Flex Consumption Plan for non-China Azure regions, ahead of the Y1 plan's retirement on September 30, 2028.

— Clear before/after and deadline, but no user action required

Agentless Vulnerability & Threat Detect…

THINNER COVERAGE BELOW

04 **HTTPS-only enforcement for Agentless V&TD Function Apps**

50

BREAKING

Enforces `https_only` on Azure Agentless Vulnerability & Threat Detection and VNet Flow Logs Function Apps to strengthen transport security.

— Names the exact config flag but not migration steps

Agentless Vulnerability & Threat Detect…

05 **ZTSA module updates for Android/ChromeOS** IMPROVED 45

The Zero Trust Secure Access Module for Android/ChromeOS (v2.0.0.1651) adds support for using an MDM-provided phone number as fallback device information, adds policy support for managing notification and storage permissions on devices, and improves ZTNA session stability/connection resilience along with PAC and SWG traffic routing.

— Lists what changed but no config, flag, or metric given

Zero Trust Secure Access Module for And…

06 **Expanded Azure region coverage for Agentless V&TD** NEW

15

Expands region support for Agentless Vulnerability & Threat Detection and VNet Flow Logs on Azure subscriptions beyond the newly added China regions.

— No specific regions or scope named Agentless Vulnerability & Threat Detect…

Scan streaming, cookie-authenticated, and plain-text custom AI endpoints in AI Scanner

BLOG ↗ STABLE

AI Scanner adds streaming, session-cookie auth, and plain-text response support for custom AI endpoints

└──▷ **USE IT**

Scan a cookie-authenticated, streaming chat endpoint that returns plain-text chunks — combining all three new capabilities in one custom target config.

yaml

```yaml
auth_type: cookie
api_key_env: SESSION_COOKIE
stream: true
response_format: text
response_template: '{{response}}'
```

Scan a streaming JSON endpoint (e.g. an OpenAI-compatible API) by enabling the `stream` option and pointing the delta path at the incremental content field.

yaml

```yaml
stream: true
response_format: json
response_template: 'choices[0].delta.content'
```

Adds `stream` option for custom AI Scanner targets to scan endpoints that return incremental responses (SSE or streamed text chunks), enabling coverage of chat-style APIs.

In JSON streaming mode, AI Scanner applies an existing response template to each SSE/streamed frame as a delta path (e.g. `choices[0].delta.content`) and concatenates extracted content in arrival order.

In text streaming mode, AI Scanner reads raw response bytes and joins them until the stream ends.

Adds `cookie` as a supported value for `auth_type`, allowing AI Scanner to authenticate against endpoints that require a browser or application session; credentials are read from `api_key_env` and sent in the `cookie` request header.

Adds `text` as a supported value for `response_format`, causing AI Scanner to treat the entire response body as raw text; response templates must include the `{{response}}` marker to insert the returned text into the response envelope.

+1 more

`Zero Trust Secure Access Module for Android/ChromeOS version 2.0.0.1651 (2026-09-28)`

BLOG ↗ LATEST STABLE

ZTSA Module for Android/ChromeOS 2.0.0.1651 adds MDM phone-number fallback and new permission policy controls.

Adds support for using an MDM-provided phone number as fallback device information.

Adds policy support for managing notification and storage permissions on Android/ChromeOS devices.

Improves ZTNA session stability and connection resilience.

Improves PAC and SWG traffic routing.

`Agentless Vulnerability & Threat Detection version 1.206.09221355 (2026-09-25)`

BLOG ↗ STABLE

Agentless V&TD for Azure gains Flex Consumption Plan migration, China region support, and expanded regional coverage.

Migrates Agentless Vulnerability & Threat Detection and VNet Flow Logs from the Y1 Consumption Plan to the Flex Consumption Plan for non-China Azure regions, ahead of the Y1 plan retirement on September 30, 2028.

Adds China region support (China East 2, China North 2, China North 3) for Agentless Vulnerability & Threat Detection and VNet Flow Logs, including vulnerability and malware scanning for Azure Virtual Machines, Disks, Azure Container Registry images, and Azure Function Apps, plus VNet Flow Logs collection and forwarding to TrendAI Vision One.

Expands region support for Agentless Vulnerability & Threat Detection and VNet Flow Logs on Azure subscriptions.

Enforces `https_only` on Azure Agentless Vulnerability & Threat Detection and VNet Flow Logs Function Apps to strengthen transport security.

WAS THIS USEFUL?  

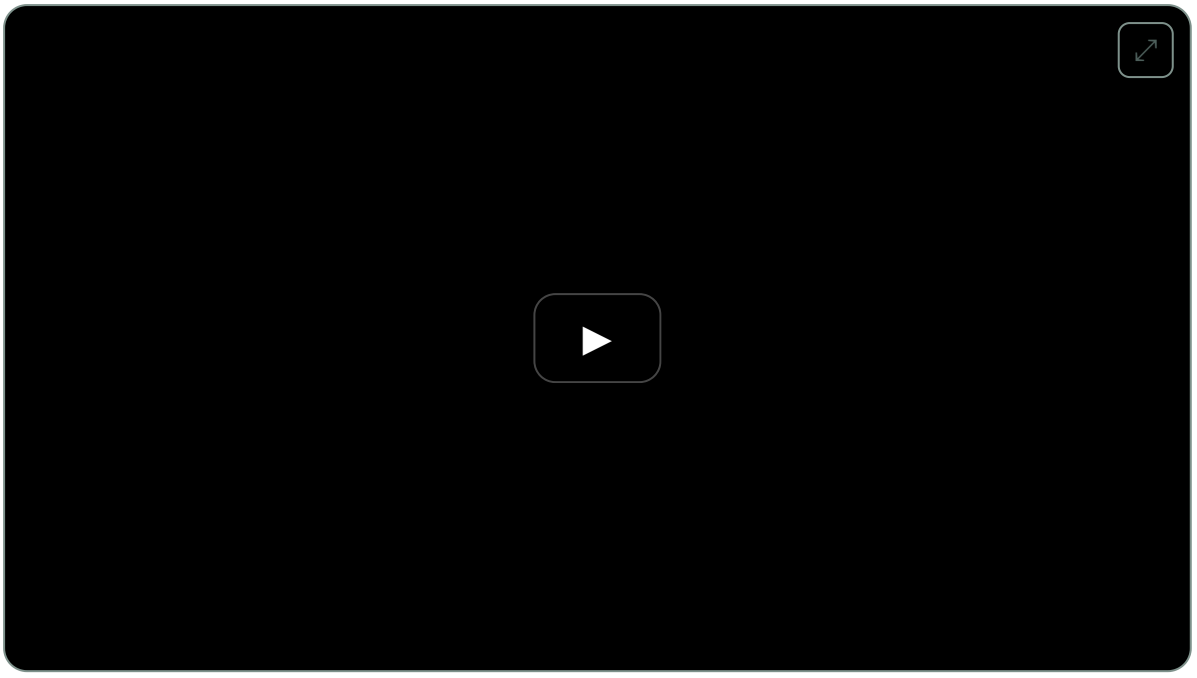
Iru

Feedback ⓘ What's New in Iru | September 25th NOTES ↗ STABLE

RANK

Iru is a device-centric platform that unifies security, IT, and compliance around devices and everything they touch.

Iru adds MSP Billing, LAPS for Mac, compliance evidence enhancements, and macOS 27 support.



Adds MSP Billing with itemized billing to increase visibility into MSP margins.

Adds LAPS for Mac as a Library Item.

Enhances Iru Compliance evidence collection.

NEW COMPLIANCE FRAMEWORK Windows Managed OS patch compliance status reporting

▾ 4 more releases in this issue · 2026-09-18 → 2026-09-25

> Recommended privacy permissions for iPhone and iPad NOTES ↗ STABLE
>
> Iru's Privacy Library Item gains a Recommended privacy permissions area for iPhone and iPad, letting admins bundle app and Safari permission prompts into a single consent scr



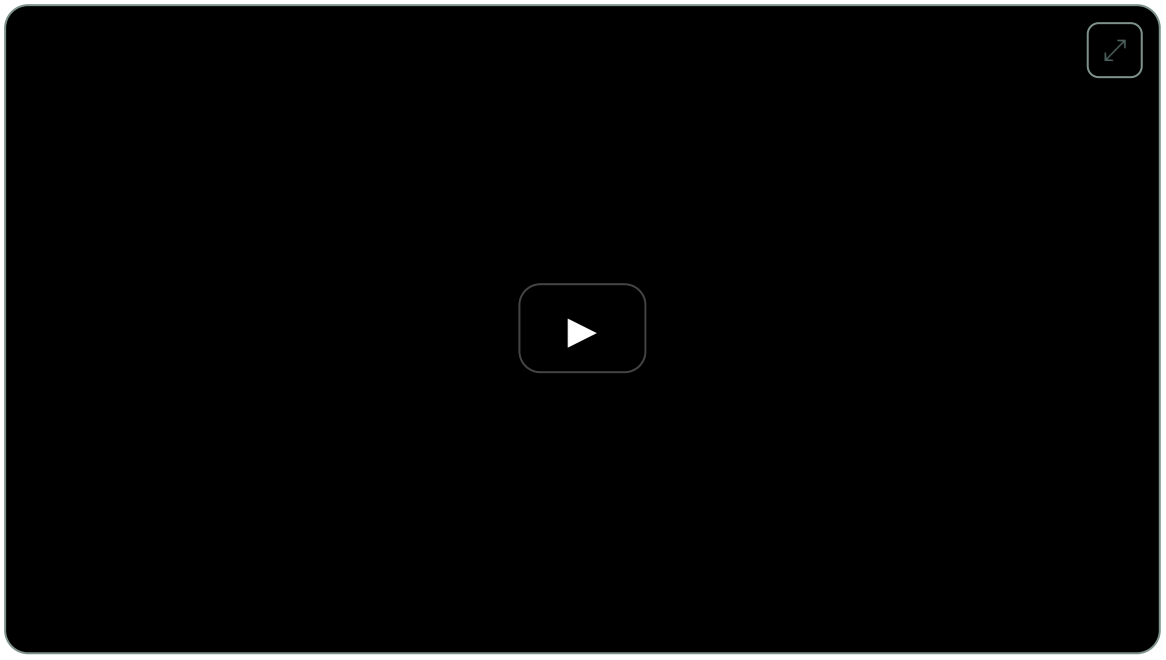
Introducing MSP Billing NOTES ↗ STABLE

Iru adds MSP Billing with daily per-tenant spend visibility and self-serve setup in the Partner Portal.

.

What's New in Iru | September 18th NOTES ↗ STABLE

Iru adds Windows Vulnerability Management, OS update status visibility, and enrollment invite email routing for Identity.



Adds Vulnerability Management and Response deployment support for Windows devices.

Adds OS update status visibility per Windows device, showing clear update state across your fleet.

Enables control over which email address enrollment invites are sent to in Iru Identity.

Windows Managed OS patch compliance NOTES ↗ STABLE

Windows Managed OS Library Items now surface per-device patch compliance with Pass/Update past due/Error status and build-level enforcement date detail.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.260928.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI ships XTM One-based Enterprise Edition licensing with a new Filigran Experience settings page, alongside SSVC vulnerability prioritization aligned to CISA BOD 26-04, Indicator-to-Security-Platform sightings, always-on ingestion logs, investigation duplication, and per-node Redis TLS validation.

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

01 XTM One Enterprise Edition licensing NEW

90

Adds `xtm:xtm_one_url` (env `XTM__XTM_ONE_URL`, default `https://common.xtm1.filigran.io`) and `xtm:xtm_one_token` (env `XTM__XTM_ONE_TOKEN`) config keys so OpenCTI registers with XTM One every 5 minutes and receives a Filigran-signed Enterprise Edition license; OpenCTI now verifies that signed license and grants EE access from it. A new `Settings > Filigran Experience` page shows the license source (OpenCTI or XTM One), customer, type, and validity dates, rechecked every 5 minutes, with a 90-day grace period for expired licenses (no grace for `trial` or `ci` types). Requires XTM One release #3831 or later — against an older XTM One the license grants nothing and OpenCTI logs a warning.

Enable XTM One-based Enterprise Edition licensing so OpenCTI automatically receives and renews its EE license from the XTM One server every 5 minutes.

yaml

```yaml
xtm:
  xtm_one_url: "https://common.xtm1.filigran.io"
  xtm_one_token: "<your-xtm-one-registration-token>"
```

— names config keys, env vars, defaults, UI path across two releases 7.260928.0

02 Per-node Redis TLS validation NEW

83

Adds `redis:tls_cluster_node_mode` (env `REDIS__TLS_CLUSTER_NODE_MODE`) config key that, when set, validates each Redis cluster node certificate individually rather than globally.

Enforce individual TLS certificate validation for each Redis cluster node instead of a single shared check.

```yaml
redis:
  tls_cluster_node_mode: true
```

— exact config key, env var, and runnable example given `product docs`

THINNER COVERAGE BELOW

03 **SSVC vulnerability prioritization** NEW 40

Adds SSVC risk-based prioritization metrics aligned to CISA BOD 26-04 for vulnerability objects.

— names the standard but not scoring mechanism or UI `7.260928.0`

04 **Observed Data occurrence counters** NEW 40

Adds `number_seen` and `max_distinct_count` counters to the Observed Data data model.

— names the fields but not how they populate or display `7.260928.0`

05 **Ingestion logs for all ingestion types** IMPROVED 40

Adds ingestion logs across all ingestion types; the previous feature flag has been removed so logging is now always active.

— notes flag removal but not log location or format `7.260928.0`

06 **Indicator-to-Security-Platform sightings** NEW 35

Supports sighting relationships between an Indicator and a Security Platform entity.

— names entity types but no workflow detail `7.260928.0`

07 **Investigation duplication** NEW 20

Supports duplicating an investigation.

— bare feature name with no mechanism or UI path `7.260928.0`

└──▷ **BREAKING ON UPGRADE**

! XTM One-based Enterprise Edition licensing requires XTM One release #3831 or later; against an older XTM One the XTM license grants nothing and OpenCTI logs a warning — upgrade XTM One before OpenCTI.

`7.260928.0` NOTES ↗ LATEST STABLE

OpenCTI 7.260928.0 adds SSVC vulnerability prioritization, Indicator-to-Security-Platform sightings, ingestion logs, and investigation duplication.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260928.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260928.0
```

Adds `number_seen` and `max_distinct_count` counters to the Observed Data data model.

Adds ingestion logs across all ingestion types (feature flag removed, now always active).

Adds SSVC risk-based prioritization metrics aligned to CISA BOD 26-04 for vulnerability objects.

Supports `sighting` relationships between an Indicator and a `Security Platform` entity.

Supports duplicating an investigation.

+1 more

WAS THIS USEFUL? 👍 👎

MISP

Feedback ⓘ v2.5.48 NOTES ↗ RANK

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

> MISP v2.5.48 adds XLSX as a native export format across events, attributes and sightings, and lets enrichment/import modules remove tags from existing attributes instead of only adding new ones.

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

01 XLSX export for events, attributes, sightings NEW

88

MISP adds XLSX export for events, attributes, and sightings via `restSearch` and the per-event and mass export pickers. A new `escape_formulas_literal:1` parameter switches formula-safe cells to literal apostrophe form instead of using `quotePrefix` cell styling, avoiding Excel formula-injection warnings.

Export all attributes matching a network indicator search as an Excel workbook for analyst review — avoiding CSV formula-injection warnings in Excel.

```
$ curl -s -H 'Authorization: <api-key>' -H 'Accept:
application/vnd.openxmlformats-
officedocument.spreadsheetml.sheet' 'https://<misp-
host>/attributes/restSearch/returnFormat:xlsx/type:ip-dst' -o
results.xlsx
```

Export attributes as XLSX with formula cells written as literal apostrophe-prefixed strings, for workflows that re-export the workbook back to CSV without preserving `quotePrefix` styling.

```
$ curl -s -H 'Authorization: <api-key>' -H 'Accept:
application/vnd.openxmlformats-
officedocument.spreadsheetml.sheet' 'https://<misp-
host>/attributes/restSearch/returnFormat:xlsx/escape_formulas_lit
eral:1' -o results_literal.xlsx
```

— Names endpoint params and exact flag with runnable curl examples v2.5.48

02 Tag removal from module results NEW

62

Enrichment and import modules can now include a `remove_tags` list in result rows, letting modules strip tags from existing attributes instead of creating duplicates.

Removed tags are surfaced in the resolution screen and in `EventShell` enrichment output.

— Names config field and output surface but no runnable example v2.5.48

v2.5.48 NOTES ↗ LATEST STABLE

MISP v2.5.48 adds XLSX export for events/attributes/sightings and tag removal from misp-module enrichment results.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.5.48 https://github.com/MISP/MISP.git
# already have the repo? check out this version:
$ git checkout v2.5.48
```

└──▷ **TRY IT**

Export all attributes matching a network indicator search as an Excel workbook for analyst review — avoiding CSV formula-injection warnings in Excel.

```
$ curl -s -H 'Authorization: <api-key>' -H 'Accept:
application/vnd.openxmlformats-officedocument.spreadsheetml.sheet'
'https://<misp-
host>/attributes/restSearch/returnFormat:xlsx/type:ip-dst' -o
results.xlsx
```

Export attributes as XLSX with formula cells written as literal apostrophe-prefixed strings, for workflows that re-export the workbook back to CSV without preserving `quotePrefix` styling.

```
$ curl -s -H 'Authorization: <api-key>' -H 'Accept:
application/vnd.openxmlformats-officedocument.spreadsheetml.sheet'
'https://<misp-
host>/attributes/restSearch/returnFormat:xlsx/escape_formulas_litera
l:1' -o results_literal.xlsx
```

Adds XLSX export for events, attributes, and sightings via `restSearch` and the per-event and mass export pickers, with `escape_formulas_literal:1` to switch formula-safe cells to literal apostrophe form instead of `quotePrefix` cell style.

Adds `remove_tags` list support in enrichment and import module result rows, letting modules strip tags from existing attributes instead of creating duplicates,

> with removed tags surfaced in the resolution screen and `EventShell` enrichment output.

WAS THIS USEFUL? 👍 👎

Recorded Future

Feedback ⓘ announced · Sep 28, 2026 BLOG ↗

RANK

Get started with the Recorded Future API. Covers authentication, base URLs, API key setup, and how to integrate threat intelligence into your security workflows.

> Recorded Future launched its MCP (Model Context Protocol) gateway to general availability, giving OAuth-authenticated AI agents access to 80+ intelligence tools, write-back capability, and direct Intelligence Graph traversal.

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

01 **Recorded Future MCP gateway reaches GA** NEW 90

Recorded Future MCP is a Model Context Protocol gateway exposing 80+ intelligence tools — including threat actor profiles, Risk Scores, ransomware metadata, malware sandbox data, and dark web intelligence — to AI agents via `OAuth 2.0` (Client ID/Secret). It is accessible as a tile in the Integration Center, supports both `OAuth 2.0` Client ID/Secret (for integration users) and portal user-based authentication (SSO), and is compatible with Claude, ChatGPT Enterprise, Copilot, Cursor, Gemini CLI, and any other OAuth-capable MCP client.

Activate Recorded Future MCP from within the platform so your AI agents gain authenticated access to the full intelligence tool catalog.

> 📍 In the Recorded Future Platform, go to Integration Center › find the Recorded Future MCP tile › click to activate › configure an integration user with Client ID/Secret (OAuth 2.0) for persistent agent access.

— Names auth method, activation path, and client compatibility with exact steps.

launch-20260928-306293b1

THINNER COVERAGE BELOW

02 **Intelligence Graph traversal for MCP agents** NEW 50

MCP supports direct graph traversal of the Intelligence Graph, letting agents find related entities and context in fewer calls, reducing token cost and latency at agentic scale.

— Explains the benefit and mechanism but no concrete usage example.

launch-20260928-306293b1

03 **Agent write-back to Watch Lists and Analyst Notes** NEW 45

MCP-connected agents can now write back to the Recorded Future Platform, updating Watch Lists and authoring Platform Analyst Notes directly rather than only reading intelligence.

— States the capability but no mechanism or steps to invoke it. launch-20260928-306293b1

announced · Sep 28, 2026 BLOG ↗

Recorded Future MCP is now GA, giving AI agents OAuth-authenticated access to 80+ intelligence tools including write-back to Watch Lists and Analyst Notes.

└──▷ HOW TO FIND IT

Activate Recorded Future MCP from within the platform so your AI agents gain authenticated access to the full intelligence tool catalog.

📍 In the Recorded Future Platform, go to Integration Center › find the Recorded Future MCP tile › click to activate › configure an integration user with Client ID/Secret (OAuth 2.0) for persistent agent access.

Adds Recorded Future MCP, a Model Context Protocol gateway exposing 80+ intelligence tools — including threat actor profiles, Risk Scores, ransomware metadata, malware sandbox data, and dark web intelligence — to AI agents via OAuth 2.0 (Client ID/Secret).

Enables write-back from agents to Recorded Future Platform: agents can update Watch Lists and author Platform Analyst Notes directly, not just read intelligence.

Accessible as a tile in the Integration Center for activation alongside existing integrations; supports both OAuth 2.0 Client ID/Secret (for integration users) and portal user-based authentication (SSO).

Supports direct graph traversal of the Intelligence Graph, letting agents find related entities and context in fewer calls to reduce token cost and latency at agentic scale.

Compatible with Claude, ChatGPT Enterprise, Copilot, Cursor, Gemini CLI, and any other OAuth-capable MCP client — no custom agent required.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Shuffle

Feedback ⓘ v2.3.0 NOTES ↗ RANK

Shuffle: A general purpose security automation platform. Our focus is on collaboration and resource sharing.

> Shuffle v2.3.0 focuses on self-hosted and air-gapped deployments, adding on-prem/hybrid AI Agents, private-registry and air-gapped app builds, and automatic recovery of stuck workflow executions.

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

01 **Air-gapped and private-registry app builds** NEW 85

Adds `SHUFFLE_BASE_IMAGE_REGISTRY` and `SHUFFLE_BASE_IMAGE_NAME` config settings for consistent custom/private registry behavior across all deployment types, and supports a custom App SDK image for building apps in disconnected or air-gapped environments on self-hosted installations.

Point a self-hosted Shuffle deployment at a private container registry so all app builds pull from your internal mirror instead of Docker Hub.

ini

```
SHUFFLE_BASE_IMAGE_REGISTRY=registry.internal.corp
SHUFFLE_BASE_IMAGE_NAME=shuffle/shuffle-app-sdk
```

— Names two env vars with a runnable config example v2.3.0

02 **On-prem and hybrid AI Agents** NEW 60

AI Agents can now run in on-premises and hybrid setups, picking the correct backend address and tool timeouts, recovering late-arriving results, and using local AI credentials or Shuffle AI synced from the cloud.

— Explains mechanism but no config surface given v2.3.0

THINNER COVERAGE BELOW

03 **Shuffle Security in Compose and Helm deployments** NEW 50

Includes Shuffle Security in both the Docker Compose configuration and Kubernetes Helm chart for easier self-hosted deployment.

— Names deployment files but no exact keys given v2.3.0

04 **Automatic rerun of stuck executions in Orborus** NEW

45

Adds an automatic rerun system within Orborus that detects and restarts stuck workflow executions, reducing the need for manual intervention.

— Explains behavior, no flag or config named v2.3.0

05 **Live workflow update streaming improvements** IMPROVED

45

Improves live workflow update streaming, including AI-generated edits, with more reliable change history and reconnection handling, and reduces background polling when only one person is viewing a workflow.

— Behavior described, no named surface or metric v2.3.0

06 **AI Agent and MCP action name matching refinement** IMPROVED

45

Refines AI Agent and MCP action app/tool name matching by searching workflows to find the correct action, with added execution and permission checks on agent-initiated runs.

— Describes mechanism, no named endpoint or flag v2.3.0

v2.3.0 NOTES ↗ LATEST STABLE

Shuffle v2.3.0 brings on-prem/hybrid AI Agents, air-gapped app builds, and automatic stuck-execution reruns via Orborus.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.3.0 https://github.com/Shuffle/Shuffle.git
# already have the repo? check out this version:
$ git checkout v2.3.0
```

└──▷ USE IT

Point a self-hosted Shuffle deployment at a private container registry so all app builds pull from your internal mirror instead of Docker Hub.

.env

```
SHUFFLE_BASE_IMAGE_REGISTRY=registry.internal.corp
SHUFFLE_BASE_IMAGE_NAME=shuffle/shuffle-app-sdk
```

Adds `SHUFFLE_BASE_IMAGE_REGISTRY` and `SHUFFLE_BASE_IMAGE_NAME` config settings for consistent custom/private registry behavior across all deployment types.

Enables AI Agents to run in on-premises and hybrid setups, picking the correct backend address and tool timeouts, recovering late-arriving results, and using local

AI credentials or Shuffle AI synced from the cloud.

Adds automatic rerun system within Orborus to detect and restart stuck workflow executions, reducing the need for manual intervention.

Supports a custom App SDK image for building apps in disconnected or air-gapped environments on self-hosted installations.

Improves live workflow update streaming, including AI-generated edits, with more reliable change history and reconnection handling, and reduces background polling when only one person is viewing a workflow.

+2 more

WAS THIS USEFUL?  

◆ Security Data Pipeline

Cribl

Feedback ⓘ changelog · Sep 29, 2026 DOCS ↗ BLOG ↗ RANK

Cribl is built for IT and Security data and provides a unified data management platform for exploring, collecting, processing, and accessing that data at scale.

> Cribl launched two major new products at CriblCon 26 — Cribl Detect, a full SIEM with 8,000+ detection rules and posture management, and StreamAI, an enterprise AI gateway with model routing and spend/redaction controls — alongside a new agentic development environment, expanded MCP/SDK agent access, AI auto-parsing, and a growing App Framework marketplace.

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

01 **Cribl Detect SIEM launch** NEW 83

Cribl Detect is a full SIEM built on the Cribl platform, combining detection, Detection Posture Management (continuously showing coverage gaps, broken rules, and missing telemetry required to detect specific attacks), investigation, response, and threat intelligence. It ships with a catalog of more than 8,000 detection rules with correlation across signals to surface multi-stage attacks, and generates detection signals directly in the data pipeline for real-time threat identification while telemetry is still in motion. It is available now for Cribl.Cloud customers at up to 50% less cost than incumbent SIEMs, priced on infrastructure rather than per-GB ingestion.

Get started with Cribl Detect on Cribl.Cloud to immediately access SIEM detection, posture management, and 8,000+ detection rules without migrating existing telemetry pipelines.

> 📍 In Cribl.Cloud, log in to your organization, navigate to the Cribl Detect tile, and select Get Started to activate detection on your existing telemetry streams.

— Mechanism, coverage and pricing detailed; only a UI-path example given

2026-09-29-1d8f4bc7

02 **StreamAI enterprise AI gateway with model routing** NEW 65

StreamAI is an enterprise AI gateway with a built-in model router that routes each prompt to the right model based on SecIT Bench benchmarks, offering free inference for customers using automatic routing. It enforces spending budgets and cutoffs, redacts sensitive data in both prompts and responses, and records every model call as audit-ready telemetry.

— Names controls and benchmark but no config keys or endpoints 2026-09-29-1d8f4bc7

THINNER COVERAGE BELOW

03 **Agentic development environment for apps and workflows** NEW 48

A new agentic development environment lets teams build apps, workflows, and agents using natural language, running on the same agentic runtime that powers Cribl Detect, with compute, scheduling, logging, and permissions included.

— Describes runtime components but no concrete usage path 2026-09-29-1d8f4bc7

04 **Cribl App Framework grows to 150+ apps** IMPROVED 40

The Cribl App Framework grows to more than 150 apps from Cribl, partners, and customers, available in the Cribl Marketplace.

— States count and location but no usage instructions 2026-09-29-1d8f4bc7

05 **Agent access via MCP and SDKs** NEW 35

Cribl opens the platform to agents via MCP and SDKs, adding 70 new integrations this year.

— Names MCP/SDK surfaces but gives no usage specifics 2026-09-29-1d8f4bc7

06 **AI auto-parsing of incoming telemetry** NEW 28

Adds AI auto-parsing of telemetry so incoming data is structured and usable the moment it arrives.

— Bare description with no mechanism or config detail 2026-09-29-1d8f4bc7

changelog · Sep 29, 2026 DOCS ↗ BLOG ↗ LATEST STABLE

Cribl launches StreamAI (AI gateway with model routing) and Cribl Detect (full SIEM with 8,000+ detection rules) at CriblCon 26.

└──▷ **HOW TO FIND IT**

Get started with Cribl Detect on Cribl.Cloud to immediately access SIEM detection, posture management, and 8,000+ detection rules without migrating existing telemetry pipelines.

📍 In Cribl.Cloud, log in to your organization, navigate to the Cribl Detect tile, and select Get Started to activate detection on your existing telemetry streams.

Adds Cribl Detect, a complete SIEM built on the Cribl platform with detection, coverage and posture management, investigation, response, and threat intelligence — available now for Cribl.Cloud customers at up to 50% less cost than incumbent SIEMs, priced on infrastructure rather than per-GB ingestion.

Adds StreamAI, an enterprise AI gateway with a built-in model router that routes each prompt to the right model based on SecIT Bench benchmarks, with free

inference for customers using automatic routing.

StreamAI enforces spending budgets and cutoffs, redacts sensitive data in both prompts and responses, and records every model call as audit-ready telemetry.

Cribl Detect includes Detection Posture Management, continuously showing coverage gaps, broken rules, and missing telemetry required to detect specific attacks.

Cribl Detect ships with a catalog of more than 8,000 detection rules covering a broad range of threats, with correlation across signals to surface multi-stage attacks.

+5 more

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Rapid7 InsightVM

Feedback ⓘ changelog · Sep 29, 2026 DOCS ↗

BLOG ↗ RANK

Rapid7 InsightVM is a vulnerability management platform that discovers, prioritizes, and remediates security risks across IT infrastructure.

The Rapid7 Agent's Windows filesystem walk tool is improved to reduce assessment failures on large filesystems and remove prior file search scope limits.

NEW SUPPORTED PLATFORM Microsoft 365 Copilot for Desktop fingerprinting

NEW SUPPORTED PLATFORM Microsoft Office LTSC 2024 fingerprinting

UPDATED SUPPORTED PLATFORM Microsoft Edge Stable and Extended Stable channel fingerprinting

NEW SUPPORTED PLATFORM Microsoft Teams fingerprinting

UPDATED SUPPORTED PLATFORM Microsoft .NET Framework 1.1 fingerprinting

WAS THIS USEFUL? 👍 👎

Future Architect Vuls

Feedback ⓘ v0.41.0-rc.2 NOTES ↗

CODE ↗ RANK

Agent-less vulnerability scanner for Linux, FreeBSD, Container, WordPress, Programming language libraries, Network devices

> Vuls v0.41.0-rc.2 expands vulnerability detection coverage through its vuls2 engine, adding support for Oracle Solaris, Apple security releases, and OpenSSH advisories.

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

01 **vuls2 detection engine adds Solaris, Apple, and OpenSSH coverage**

NEW 35

The vuls2 detection engine now detects vulnerabilities in the Oracle Solaris family, Apple security releases, and OpenSSH security advisories.

— Names three new detection targets but no flags, commands, or mechanism detail.

v0.41.0-rc.2

v0.41.0-rc.2 NOTES ↗ CODE ↗ PRE-RELEASE

Vuls v0.41.0-rc.2 adds Oracle Solaris, Apple security releases, and OpenSSH advisory detection via vuls2

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.41.0-rc.2 https://github.com/future-architect/vuls.git
# already have the repo? check out this version:
$ git checkout v0.41.0-rc.2
```

Detects vulnerabilities in the Oracle Solaris family through the vuls2 detection engine.

Detects Apple security releases through the vuls2 detection engine.

Detects OpenSSH security advisories through the vuls2 detection engine.

WAS THIS USEFUL? 👍 👎

Qualys Patch Management

Feedback ⓘ

changelog · Sep 29, 2026 | BLOG ↗ | RANK

Qualys Patch Management is a cloud-based vulnerability management tool that identifies, prioritizes, and automates the deployment of security patches across IT infrastructure.

> Qualys TruRisk Eliminate now lets teams block risky patches environment-wide via QQL-based rules and automatically skip low-reliability patches in zero-touch jobs using AI-driven reliability scoring, backed by curated interim mitigations.

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

01 Patch Blocking Rules in TruRisk Eliminate NEW 95

Adds Patch Blocking Rules under `TE > Configuration > Patch Blocking Rules`, available for Windows, Linux, and Mac in release 4.2, letting admins define patches to block via a QQL query (e.g. `patch.kb: [KB5002907, KB5124008, KB5124012, KB5123099]`) so they are excluded across every deployment job at once without editing individual jobs. Scope is controlled by up to five asset tags per rule, with up to 10 rules allowed per platform, and each rule requires a name and a blocking reason.

Immediately stop a vendor-paused or known-bad Windows patch from deploying across your entire environment without touching individual patch jobs.

📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. Block paused Sep-2026 updates) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907, KB5124008, KB5124012, KB5123099] and click Preview to confirm the match. Leave asset tags empty to cover all assets, then save.

Protect production assets while a test group validates a new patch by scoping a blocking rule to production asset tags.

📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. Hold KB5002907 from prod during validation) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907] and click Preview. Under asset scope, add up to five production asset tags to restrict the block to production only, then save.

— Names exact UI path, QQL syntax, and rule limits. 2026-09-29-7aca09ff

02 AI-Powered Patch Reliability Scoring NEW 70

Adds AI-Powered Patch Reliability Scoring, which combines LLM analysis of real-world internet feedback with Qualys rollback and vulnerability-reopen telemetry to rate each

patch High, Medium, Low, or Unidentified. Adding reliability intelligence to a zero-touch job's Select Patches step QQL automatically skips Low-reliability patches from deployment while other patches continue as scheduled.

Keep zero-touch patch jobs running while automatically skipping low-reliability patches by adding reliability intelligence to the job's patch-selection query.

> 📍 In the zero-touch patch job's Select Patches step, update the QQL to include reliability criteria alongside your existing filters so that Low-reliability patches — such as those currently rated Low by AI-Powered Patch Reliability Scoring — are excluded from automatic deployment while all other patches continue to deploy.

— Explains rating mechanism but omits exact QQL field syntax. 2026-09-29-7aca09ff

THINNER COVERAGE BELOW

03 **Curated mitigations for held patches** NEW 22

Introduces Qualys-curated mitigations as a complementary layer to reduce exposure while a blocked or low-reliability patch is held in testing and staging.

— Single-sentence description with no named surface or steps. 2026-09-29-7aca09ff

└──▷ ALSO FROM THESE RELEASES

changelog · Sep 29, 2026 BLOG ↗ LATEST STABLE

Qualys TruRisk Eliminate adds Patch Blocking Rules and AI-Powered Patch Reliability Scoring to keep bad patches out of automation.

└──▷ HOW TO FIND IT

Immediately stop a vendor-paused or known-bad Windows patch from deploying across your entire environment without touching individual patch jobs.

> 📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. Block paused Sep-2026 updates) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907, KB5124008, KB5124012, KB5123099] and click Preview to confirm the match. Leave asset tags empty to cover all assets, then save.

Protect production assets while a test group validates a new patch by scoping a blocking rule to production asset tags.

> 📍 In TruRisk Eliminate, go to TE › Configuration › Patch Blocking Rules, select Windows, click Create Rule. Enter a name (e.g. Hold KB5002907 from prod during validation) and a blocking reason. Under Criteria (QQL), enter: patch.kb: [KB5002907]

and click Preview. Under asset scope, add up to five production asset tags to restrict the block to production only, then save.

Keep zero-touch patch jobs running while automatically skipping low-reliability patches by adding reliability intelligence to the job's patch-selection query.

📍 In the zero-touch patch job's Select Patches step, update the QQL to include reliability criteria alongside your existing filters so that Low-reliability patches — such as those currently rated Low by AI-Powered Patch Reliability Scoring — are excluded from automatic deployment while all other patches continue to deploy.

Adds `Patch Blocking Rules` in `TE > Configuration > Patch Blocking Rules` (available for Windows, Linux, and Mac in release 4.2), letting you define patches via a QQL query and block them across every deployment job at once without editing individual jobs; scope is controlled by up to five asset tags per rule, with up to 10 rules per platform, and each rule requires a name and blocking reason.

Adds AI-Powered Patch Reliability Scoring that combines LLM analysis of real-world internet feedback with Qualys rollback and vulnerability-reopen telemetry to rate each patch High, Medium, Low, or Unidentified — automatically skipping Low-reliability patches from zero-touch jobs when reliability intelligence is added to the job's QQL in the `Select Patches` step.

Introduces Qualys-curated mitigations as a complementary layer to reduce exposure while a blocked or low-reliability patch is held in testing and staging.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Censys

Feedback ⓘ September 28, 2026 (2026-09-28) BLOG ↗ RANK

Censys is a search engine for internet-connected devices that indexes certificates, websites, and network services for security research and reconnaissance.

> Censys shipped an AI Skills plugin that lets Claude agents drive cencli directly, raised the Adversary Investigation batch limit sixfold, and added two new Kiteworks fingerprints to the Platform.

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

01 **AI Skills plugin for cencli agents** NEW 75

Adds the `censys-ai-skills` repo, a Claude plugin that exposes `cencli` commands to AI agents for asset lookups, investigations, and CensEye job execution. It can be installed as a Claude plugin and prompted to run asset lookups and investigations autonomously.

Wire cencli into a Claude agent so it can autonomously look up assets and run investigations without leaving the AI workflow.

> 📍 Install the censys-ai-skills repo as a Claude plugin, then prompt your agent: 'Use cencli to look up all assets tagged Kiteworks PCN and run an investigation on the top 100 indicators.'

— Named repo and commands plus install/usage steps September 28, 2026 (2026-09-28)

02 **Adversary Investigation batch limit raised to 100** IMPROVED 60

Raises the Adversary Investigation AI investigator batch limit from 35 to 100 indicators per submission.

— Exact before/after numbers but no usage command September 28, 2026 (2026-09-28)

THINNER COVERAGE BELOW

03 **New Platform fingerprints for Kiteworks** NEW 45

Adds two new fingerprints to the Platform: `software Kiteworks PCN` and `software Kiteworks Secure MFT Server`.

— Named fingerprints but no usage guidance given September 28, 2026 (2026-09-28)

> September 28, 2026 (2026-09-28) BLOG ↗ LATEST STABLE

Censys adds AI Skills plugin for cencli-powered agents and raises AI investigator batch limit to 100 indicators.

└──▷ **HOW TO FIND IT**

Wire cencli into a Claude agent so it can autonomously look up assets and run investigations without leaving the AI workflow.

> 📍 Install the censys-ai-skills repo as a Claude plugin, then prompt your agent: 'Use cencli to look up all assets tagged Kiteworks PCN and run an investigation on the top 100 indicators.'

Adds the `censys-ai-skills` repo, a Claude plugin that exposes `cencli` commands to AI agents for asset lookups, investigations, and CensEye job execution.

Raises the Adversary Investigation AI investigator batch limit from 35 to 100 indicators per submission.

Adds two new fingerprints to the Platform: `software Kiteworks PCN` and `software Kiteworks Secure MFT Server`.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

iDenfy

Feedback ⓘ snapshot-20260929 NOTES ↗ RANK

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

> iDenfy's mobile SDK update focuses on the Android SDK, adding landscape support, improved document selection guidance, glare/blur detection for older devices, and deeper UI customization.

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

01 **Android SDK capture and UI improvements** NEW 48

The Android SDK gains landscape mode support, a recommended document badge in the document selection UI to guide users toward the right document type, blur and glare detection for older Android devices to improve capture quality, and color scheme customization with XML fallbacks for theming the SDK's UI.

— Names four concrete SDK additions but no APIs or config keys. snapshot-20260929

snapshot-20260929 NOTES ↗ LATEST STABLE

iDenfy mobile SDKs gain landscape mode, recommended document badge, blur/glare detection, and Smart-ID refresh state.

Adds landscape mode support to the Android SDK.

Adds a recommended document badge in the Android SDK document selection UI.

Adds blur and glare detection for older Android devices.

Adds color scheme customization with XML fallbacks for the Android SDK.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2022 (2026-09-29) BLOG ↗

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 shipped a large slate of updates spanning identity, device management, patch management and password security — headlined by a Password Vault in the User Portal, a unified Patch Management dashboard, Access Requests approval workflows, and step-up authentication with AMR claims — alongside breaking changes that restrict LDAP TLS cipher suites and retire legacy API endpoints and filter operators.

└──▷ **WHAT SHIPPED · 25 FEATURES** most completely described first ? what's the number?

01 **Unified Patch Management dashboard** NEW 95

Adds a unified Patch Management dashboard under `Device Management > Patch Management` with `Overview`, `Apple`, and `Windows` tabs, severity filters, KB search, CVE detail review, CSV/JSON export of the KB list, and bulk remediation; admins can immediately remediate critical missing patches by applying the `Critical` severity filter, selecting affected devices, and choosing `Install missing critical patches`.

Immediately remediate critical missing patches across non-compliant Windows devices without waiting for the next policy window.

> 📍 In the Admin Portal, go to Device Management › Patch Management › Windows, apply the `Critical` severity filter, then select affected devices and choose `Install missing critical patches` to trigger immediate remediation.

— Names tabs, filters, export formats and a runnable remediation path

Release Notes 2026 (2026-09-29)

02 **User reactivation endpoint replaces deprecated GET call** 90

BREAKING

Adds `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` to replace the deprecated `GET api/systemusers/reactivate/{id}`, which now returns a 404 error; requests must include an `activation_email` field in the JSON body.

Reactivate a user account using the new POST endpoint after removing the deprecated GET call from scripts or automations.

```
$ curl -s -X POST
'https://console.jumpcloud.com/systemusers/reactivate/<user_id>'
-H 'x-api-key: <your-api-key>' -H 'Content-Type:
application/json' -d '{"activation_email": "user@example.com"}'
```

— Exact endpoints, method and required field with a runnable example

Release Notes 2024 (2026-09-29)

---

03 **Access Requests approval workflow** NEW 80

Adds Access Requests, letting admins configure an approval flow so users must request and receive admin approval before gaining access to a sensitive application; supports assigning approvers and automatic JIRA ticket creation via a configured webhook URL.

Configure an Access Request approval flow so users must request and receive admin approval before gaining access to a sensitive application, with automatic JIRA ticket creation.

> 📍 In the JumpCloud Admin Portal, go to Access Requests › Create Approval Flow, assign approvers, enable webhook integration, and enter your JIRA webhook URL so tickets are created automatically on each new request submission.

— Exact UI path and JIRA webhook integration described Release Notes 2025 (2026-09-29)

---

04 **Step-up authentication and audit signals for SSO** NEW 80

Adds a `sso_auth` event to logs for forced re-authentication events, giving audit visibility for compliance and security reviews. Adds `ForceAuthn='true'` (SAML) and `prompt=login` (OIDC) support so applications can trigger step-up authentication without ending the user's existing session, and includes Authentication Methods References (AMR) claims in SAML assertions and OIDC ID tokens, tagging logins with proof of factors used (e.g., password + OTP).

— Names exact SAML/OIDC parameters and log event but no exact API call

Release Notes 2025 (2026-09-29)

---

05 **LDAP TLS restricted to strong cipher suites** BREAKING 80

Restricts LDAP TLS connections to five strong cipher suites: `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and

`TLS_RSA_WITH_AES_128_GCM_SHA256`, removing weak CBC-mode and SHA-1/SHA-256 suites; older suites (including RSA or ECDHE key exchange) are no longer accepted.

— Exact cipher suite names but no migration command Release Notes 2025 (2026-09-29)

06 **Password Vault in User Portal** NEW 75

Adds a Password Vault in the User Portal, accessible from the left navigation after login, supporting credential creation, sharing, personal and shared folders, and bulk import via CSV or Excel file.

— Names UI location and import formats but no field-level detail

Release Notes 2026 (2026-09-29)

07 **macOS Ventura upgrade delay/block policies** NEW 75

Adds new macOS software management policies to delay or block upgrades to macOS 13 Ventura on managed devices; configured in Admin Portal under Device Management > Policies > New Policy by selecting the Ventura delay/block software management policy and binding it to a device group.

Block macOS 13 Ventura upgrades on managed devices while your team validates app compatibility.

📍 In the Admin Portal, go to Device Management › Policies › New Policy, select macOS and choose the Ventura delay/block software management policy, then bind it to the relevant device group.

— Exact UI navigation given for applying the policy Release Notes 2022 (2026-09-29)

08 **Global MFA Requirement for MSP admins** NEW 70

Adds a Global MFA Requirement setting (Account > Settings) that enforces MFA for all MSP admins across the tenant, closing gaps where individual admin accounts had MFA disabled.

Enforce MFA for all MSP admins across the tenant from the Multi-Tenant Portal to close gaps where individual admin accounts had MFA disabled.

📍 In the Admin Portal, go to Account › Settings, then locate Global MFA Requirement and enable the setting.

— Exact settings path given for enforcement Release Notes 2024 (2026-09-29)

09 **Dynamic Device Groups by attribute** IMPROVED 70

Supports automating device group membership via Dynamic Groups, using attribute rules such as OS = Linux and Vendor = Dell so matching devices are added automatically

without manual intervention.

Automate device group membership so newly enrolled Linux endpoints are added without manual intervention.

> 📍 In the Admin Portal, go to Device Groups › (select or create a group) › Group Details › Set as Dynamic, then add attribute rules such as OS = Linux and Vendor = Dell to automatically include matching devices.

– Exact attribute example given but unclear if newly added

Release Notes 2023 (2026-09-29)

---

10 **Password Manager Cloud Backup** NEW 70

Adds a Cloud Backup toggle for the Password Manager (Password Manager > Settings), giving users a secure recovery path if their local vault is lost.

Enable Cloud Backup for the Password Manager to give users a secure recovery path if their local vault is lost.

> 📍 In the Admin Portal, go to Password Manager › Settings and enable the Cloud Backup toggle for your organization.

– Exact settings path given, mechanism is a simple toggle Release Notes 2023 (2026-09-29)

---

11 **EU region selection for Password Manager** NEW 70

Supports selecting the EU region at first launch so credentials are stored and activated within the EU region; choose `EU` on the desktop app activation page before entering the activation code, or on the browser extension pairing page before pairing.

Configure the EU region at first launch so that credentials are stored and activated within the EU region.

> 📍 On the Password Manager desktop app activation page, select `EU` from the region selector before entering your activation code; for the browser extension, select `EU` on the pairing page before pairing.

– Exact selection steps given for desktop and browser extension

What's New: Password Manager Desktop Ap…

---

12 **Dynamic Groups filter operators renamed** BREAKING 65

Deprecates the `$le` and `$ge` filter operators in Dynamic Groups `memberQuery` API filters; replace with `$lte` and `$gte` respectively in all API requests.

– Exact operator names but no endpoint or example call Release Notes 2024 (2026-09-29)

13 **Agent deployment via Commands** IMPROVED 65

Supports deploying agents such as the SentinelOne Agent across a fleet from the JumpCloud Admin Portal's Commands feature, by creating a new command targeting an OS (macOS, Windows, or Linux), pasting the install script, and running it against a device group.

Deploy the SentinelOne Agent across your Linux fleet from the JumpCloud Admin Portal using the Commands feature — no manual distribution needed.

> 📍 In the Admin Portal, go to Commands › New Command, select the target OS (macOS, Windows, or Linux), paste the SentinelOne Agent installation script, and run it against your device group.

— Exact steps given but unclear if newly introduced Release Notes 2022 (2026-09-29)

---

14 **Dormant Application Access detection refinement** IMPROVED 60

Improves Dormant Application Access detection in `Access Risk > Configuration` to trigger only on re-access after inactivity (not first-time access), reducing noise from newly provisioned apps.

— Names the config location and behavior change precisely Release Notes 2026 (2026-09-29)

THINNER COVERAGE BELOW

15 **Android 15 and Windows 11 25H2 support in MDM** NEW 40

Adds full Android 15 support in MDM and Windows 11 25H2 support for JumpCloud Windows management components.

— Names both OS versions but no rollout or configuration detail

Release Notes 2024 (2026-09-29) Release Notes 2025 (2026-09-29)

---

16 **MFA frequency controls in Conditional Access (Preview)** NEW 40

Adds MFA Frequency Controls to Conditional Access Policies (Preview), letting admins define how often users must re-authenticate with MFA.

— No UI steps or config values given, still in preview product docs

---

17 **Mac Personalization policy updates for Dock and Lock Screen** IMPROVED 35

Extends Mac Personalization policies with updates to Dock Settings and Items and Lock Screen (Screen Saver) policies.

— Names the policy areas but no fields or steps Release Notes 2025 (2026-09-29)

18 **Multiple IP addresses in RADIUS server configuration** IMPROVED 30

Adds multiple IP address support to RADIUS server configuration.

— Bare description, no config keys or limits given Release Notes 2024 (2026-09-29)

19 **Custom banned password lists** NEW 30

Adds Custom Banned Password List support, enabling organizations to block organization-specific weak or prohibited passwords.

— Names the capability but no config path or limits product docs

20 **Microsoft 365 security group integration** NEW 30

Adds Microsoft 365 integration support for M365/Azure AD security groups.

— Only a one-line description, no setup steps Release Notes 2022 (2026-09-29)

21 **Workflows with templates** NEW 30

Adds a Workflows feature with a builder UI, including support for templates, as shown across the release's Workflows UI screens.

.

— Only illustrated via screenshots; no written mechanism given

Release Notes 2026 (2026-09-29)

22 **Password Manager Desktop item sorting by nickname** NEW 25

Adds item sorting by nicknames in the Password Manager Desktop App.

— Bare mention with no further mechanism What's New: Password Manager Desktop Ap…

23 **Trial recap summary** NEW 15

Adds a trial recap feature summarizing feature usage during a JumpCloud trial.

— Bare description with no mechanism or location Release Notes 2022 (2026-09-29)

24 **MCP Server** NEW 10

Adds an MCP Server as part of JumpCloud's 2025 platform additions.

— Named only in summary, no further detail given Release Notes 2025 (2026-09-29)

25 **AI Commands Builder** NEW 10

Adds an AI Commands Builder as part of JumpCloud's 2025 platform additions.

— Named only in summary, no further detail given Release Notes 2025 (2026-09-29)

└──▷ ALSO FROM THESE RELEASES

Subscribe to JumpCloud product and agent release note feeds so your RSS reader automatically surfaces new capability announcements.

> 📍 1. Go to the `Release Notes and Bug Fixes` page in the JumpCloud Help Center. 2. Click `Subscribe to Release Notes and Bug Fixes RSS Feed` to receive product release notes, agent release notes, and bug fix reports. 3. (Optional) Click `Subscribe to Agent Release Notes RSS Feed` for agent-only updates. 4. Open your RSS reader and add the feed URL to begin receiving updates automatically.

└──▷ BREAKING ON UPGRADE

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; replace all usages with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! The `$le` and `$ge` filter operators in Dynamic Groups `memberQuery` API filters are deprecated; replace with `$lte` and `$gte` respectively in all API requests.

! LDAP TLS connections that used CBC-mode or SHA-1/SHA-256 cipher suites (including those with RSA or ECDHE key exchange) are no longer accepted; only `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256` are permitted.

! The `.appimage` Linux build is no longer supported as of April 18, 2024 (v. 2.2.197); switch to the RPM or Debian build.

`Release Notes 2022 (2026-09-29)` `BLOG ↗` `STABLE`

JumpCloud: Adds new macOS software management policies to delay or block upgrades to macOS 13 Ventura on managed devices.

└──▷ HOW TO FIND IT

Deploy the SentinelOne Agent across your Linux fleet from the JumpCloud Admin Portal using the Commands feature — no manual distribution needed.

> 📍 In the Admin Portal, go to Commands › New Command, select the target OS (macOS, Windows, or Linux), paste the SentinelOne Agent installation script, and run it against your device group.

Block macOS 13 Ventura upgrades on managed devices while your team validates app compatibility.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy, select `macOS` and choose the Ventura delay/block software management policy, then bind it to the relevant device group.

Enforce disk-encryption-based Conditional Access to protect sensitive SSO resources on unencrypted endpoints.

> 📍 In the Admin Portal, go to Security › Conditional Access › New Policy, add a Device Condition for `Disk Encryption` set to `Enabled`, and apply the policy to the target SSO application.

Adds Microsoft 365 integration support for M365/Azure AD security groups.

Adds MDM Upcoming Certificate Expirations widget to the JumpCloud Homepage (single-org view).

Adds trial recap feature summarizing feature usage during a JumpCloud trial.

`Release Notes 2024 (2026-09-29)` `BLOG ↗` `STABLE`

JumpCloud: Adds `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` endpoint to replace the deprecated `GET api/systemusers/reactivate/{id}`.

└──▷ **TRY IT**

Reactivate a user account using the new POST endpoint after removing the deprecated GET call from scripts or automations.

```
$ curl -s -X POST
'https://console.jumpcloud.com/systemusers/reactivate/<user_id>' -H
'x-api-key: <your-api-key>' -H 'Content-Type: application/json' -d
'{"activation_email": "user@example.com"}'
```

Enforce MFA for all MSP admins across the tenant from the Multi-Tenant Portal to close gaps where individual admin accounts had MFA disabled.

> 📍 In the Admin Portal, go to Account › Settings, then locate `Global MFA Requirement` and enable the setting.

Adds full Android 15 support in MDM.

Adds multiple IP address support to RADIUS server configuration.

└──▷ **BREAKING ON UPGRADE**

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; replace all usages with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! The `$le` and `$ge` filter operators in Dynamic Groups `memberQuery` API filters are deprecated; replace with `$lte` and `$gte` respectively in all API requests.

Release Notes 2025 (2026-09-29) BLOG ↗ STABLE

JumpCloud adds Access Requests, MCP Server, AI Commands Builder, Shadow AI Dashboard, Custom Admin Roles, and major SSO/MFA enhancements in 2025.

└──▷ **HOW TO FIND IT**

Configure an Access Request approval flow so users must request and receive admin approval before gaining access to a sensitive application, with automatic JIRA ticket creation.

📍 In the JumpCloud Admin Portal, go to Access Requests › Create Approval Flow, assign approvers, enable webhook integration, and enter your JIRA webhook URL so tickets are created automatically on each new request submission.

Adds `sso_auth` event to logs for forced re-authentication events, providing audit visibility for compliance and security reviews.

Adds `ForceAuthn='true'` (SAML) and `prompt=login` (OIDC) support so applications can trigger step-up authentication without ending the user's existing session.

Includes Authentication Methods References (AMR) claims in SAML assertions and OIDC ID tokens, tagging logins with proof of factors used (e.g., password + OTP).

Extends Mac Personalization policies with updates to Dock Settings and Items and Lock Screen (Screen Saver) policies.

Adds Windows 11 25H2 support for JumpCloud Windows management components.

+1 more

└──▷ **BREAKING ON UPGRADE**

! LDAP TLS connections that used CBC-mode or SHA-1/SHA-256 cipher suites (including those with RSA or ECDHE key exchange) are no longer accepted; only `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256` are permitted.

Release Notes 2026 (2026-09-29) BLOG ↗ STABLE

JumpCloud: Adds Password Vault in the User Portal (accessible from left navigation after login), supporting credential creation, sharing, personal and shared folders, and bulk import via CSV or Excel file.

└──▷ **HOW TO FIND IT**

Upload an internal runbook to the AI Assistant knowledge base so the assistant can answer org-specific process questions.

> 📍 In the Admin Portal, navigate to AI Assistant › Knowledge Base, click Upload, select a PDF, TXT, DOCX, or MD file (up to 10 MB), and confirm. Repeat for up to 25 files per org.

Configure an AI agent to authenticate with Amazon Bedrock AgentCore, which requires credentials in the request body rather than the header.

> 📍 In the Admin Portal, open the agent, go to the Authentication tab, and set the client authentication type to `Client secret post` so the client ID and secret are sent in the request body when requesting a token.

Immediately remediate critical missing patches across non-compliant Windows devices without waiting for the next policy window.

> 📍 In the Admin Portal, go to Device Management › Patch Management › Windows, apply the `Critical` severity filter, then select affected devices and choose `Install missing critical patches` to trigger immediate remediation.

Adds a unified Patch Management dashboard under `Device Management > Patch Management` with `Overview`, `Apple`, and `Windows` tabs, severity filters, KB search, CVE detail review, CSV/JSON export of KB list, and bulk remediation.

Improves Dormant Application Access detection in `Access Risk > Configuration` to trigger only on re-access after inactivity (not first-time access), reducing noise from newly provisioned apps.

WAS THIS USEFUL? 👍 👎

Ping Identity

Feedback ⓘ 28 Sep 2026 (2026-09-29) BLOG ↗

RANK

Read docs, explore use cases, learn best practices

> Ping Identity's Authorize gateway now fails closed instead of continuing to run unhealthy when it can't retrieve policy, while the Advanced Identity Cloud admin console picks up a broad set of additions spanning SAML metadata import, email provider configuration, telephone user properties, node designer, navigation profiles and journey canvas UX.

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

01 Navigation profiles for hosted account pages NEW

85

Enables customization of the hosted account pages navigation menu for different user groups using navigation profiles, configured via Hosted Pages > Navigation Profiles by creating a profile, defining menu items, and assigning it to a target group.

Serve tailored navigation menus to different user groups on hosted account pages by assigning navigation profiles per group.

> 📍 In the Advanced Identity Cloud admin console, go to Hosted Pages › Navigation Profiles, create a profile for a user group, define the menu items, then assign the profile to the target group.

— Full UI path and steps make this immediately actionable. 28 Sep 2026 (2026-09-29)

02 Skip XML signature check on SAML metadata import NEW

80

Adds ability to skip XML signature verification when importing remote SAML 2.0 provider metadata by enabling Skip XML signature verification during import, avoiding failures when a partner's metadata lacks a valid signature.

Skip signature validation when federating with a partner whose SAML 2.0 metadata lacks a valid XML signature, avoiding import failures.

> 📍 In the admin console, go to Federation › Entity Providers › Import Remote Provider, upload the metadata file, and enable Skip XML signature verification before saving.

— Exact UI navigation and toggle name make this directly usable. 28 Sep 2026 (2026-09-29)

03 Entra ID and SMTP email provider configuration NEW

80

Enables configuration of Microsoft Entra ID or custom SMTP email providers and test-email sending from the Advanced Identity Cloud admin console, via Settings > Email Provider, selecting `Microsoft Entra ID` or `Custom SMTP`, entering connection details, and sending a test email to confirm delivery.

Configure a Microsoft Entra ID or custom SMTP email provider and verify delivery with a test email before rolling out to users.

> 📍 In the Advanced Identity Cloud admin console, go to Settings › Email Provider, select `Microsoft Entra ID` or `Custom SMTP`, fill in the connection details, then click `Send Test Email` to confirm.

— Exact console path and steps to configure and verify delivery. `28 Sep 2026 (2026-09-29)`

---

04 **Authorize gateway fails closed on policy retrieval failure**

75

`BREAKING`

Authorize gateway 1.2.2 now logs an error and shuts down when it cannot retrieve the latest authorization policy version from PingOne or the locally stored deployment configuration, replacing the previous behavior of remaining alive but unable to serve decision requests. Any monitoring or restart logic that assumed a running-but-unhealthy gateway process must be updated.

— Clear before/after behavior described, but no config key or command to act on.

`product docs`

---

05 **Telephone number as user property format** NEW

60

Supports telephone number as a user property format, with optional server-side validation and separate country code and phone number fields in journeys.

— Mechanism described but no exact configuration path given. `28 Sep 2026 (2026-09-29)`

---

06 **LexisNexis DDP Profiler script injection location** NEW

60

Allows configuring the LexisNexis DDP Profiler node to inject scripts in the `head` element or in an iframe.

— Names the concrete injection targets but no config steps. `28 Sep 2026 (2026-09-29)`

THINNER COVERAGE BELOW

---

07 **Target phase selection for reconciliations** NEW

50

Adds `runTargetPhase` selection in the Advanced Identity Cloud admin console when running reconciliations.

— Names the property but gives no navigation path or example. `28 Sep 2026 (2026-09-29)`

---

08 **Journey canvas panning and multi-node selection** IMPROVED

50

Enables panning journey canvases by holding the space bar and dragging with the left mouse button, and enables selecting multiple journey nodes by holding Shift and clicking each node.

— Exact key and mouse combinations given, directly usable. 28 Sep 2026 (2026-09-29)

09 **Multi-value custom node properties** NEW

45

Allows selection of multiple values from predefined options when configuring custom node properties in the node designer.

— States the capability but no steps or example to follow. 28 Sep 2026 (2026-09-29)

10 **Next-generation scripting engine now default** IMPROVED

40

Makes the next-generation scripting engine the default and recommended option when creating scripts in the Advanced Identity Cloud admin console.

— Describes the change but no detail on what differs in the engine.

28 Sep 2026 (2026-09-29)

11 **Search and sort for AI agent identities** NEW

25

Adds search and sort capabilities for the list of AI agent identities.

— Bare mention with no mechanism or navigation detail. 28 Sep 2026 (2026-09-29)

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! Authorize gateway 1.2.2 will shut down (instead of staying alive) when it cannot reach PingOne or load its local deployment configuration — any monitoring or restart logic that assumed a running-but-unhealthy gateway process must be updated.

! Authorize gateway instances that previously stayed alive in an unhealthy state when policy retrieval failed will now shut down on upgrade to 1.2.2 — any automation or health-check logic that relied on the instance remaining running in that scenario will need to be updated.

28 Sep 2026 (2026-09-29) BLOG ↗ LATEST STABLE

Ping Identity adds multi-value node properties, telephone format, Entra ID/SMTP email config, SAML signature skip, and more in Sep 2026 release.

└──▷ HOW TO FIND IT

Skip signature validation when federating with a partner whose SAML 2.0 metadata lacks a valid XML signature, avoiding import failures.

📍 In the admin console, go to Federation › Entity Providers › Import Remote Provider, upload the metadata file, and enable `Skip XML signature verification` before saving.

Configure a Microsoft Entra ID or custom SMTP email provider and verify delivery with a test email before rolling out to users.

📍 In the Advanced Identity Cloud admin console, go to Settings › Email Provider, select `Microsoft Entra ID` or `Custom SMTP`, fill in the connection details, then click `Send Test Email` to confirm.

Serve tailored navigation menus to different user groups on hosted account pages by assigning navigation profiles per group.

📍 In the Advanced Identity Cloud admin console, go to Hosted Pages › Navigation Profiles, create a profile for a user group, define the menu items, then assign the profile to the target group.

Adds `runTargetPhase` selection in the Advanced Identity Cloud admin console when running reconciliations.

Adds ability to skip XML signature verification when importing remote SAML 2.0 provider metadata.

Enables configuration of Microsoft Entra ID or custom SMTP email providers and test-email sending from the Advanced Identity Cloud admin console.

Supports telephone number as a user property format, with optional server-side validation and separate country code and phone number fields in journeys.

Allows selection of multiple values from predefined options when configuring custom node properties in the node designer.

`+6 more`

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ announced · Sep 29, 2026 BLOG ↗ RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

ConductorOne now handles OIDC sign-in, AuthZEN real-time authorization, and governed access requests for AI-built apps, eliminating local user tables and stale permission snapshots.

NEW INTEGRATION Acts as a standard OIDC provider for AI-built applications, federating authentication to the company's existing identity provider without requiring a separate user table or login page.

NEW INTEGRATION Adds an OpenID AuthZEN authorization endpoint so AI-built apps can query ConductorOne for real-time, per-request access decisions instead of relying on stale local role snapshots.

Surfaces access-denied events as governable access requests: employees can submit a request from within the application, owners review it with context, and time-bound access is granted or denied without a code change or redeploy.

Makes AI-built application access a ConductorOne entitlement, bringing it into existing request, approval, access-review, certification, and revocation workflows alongside all other governed systems.

▾ **1 more release in this issue** · 2026-09-28 → 2026-09-29

> announced · Sep 28, 2026 BLOG ↗
>
> ConductorOne launches C1 AppHub, a self-hosted internal app platform connecting AI-built apps to identity, authz, credentials, and model access.
>
> .
> .
> .
>
> Introduces C1 AppHub, a self-hosted internal application platform that wires each new app to ConductorOne identity, authorization, credential/token vending, network access, and model access without per-app custom integrations.
>
> AppHub apps are added to C1.ai's identity graph alongside every other connected system, each with an assigned owner, enabling inclusion in access reviews and offboarding workflows.
>
> Supports a skills file that teaches coding agents the company's approved development pattern — directing them to use C1.ai for sign-in, authorization, credential and token vending, network access, and model access as the app is built.
>
> Available as a public repository so teams can inspect the architecture, run it in their own environment, and adapt the integration pattern to their needs.

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ announced · Sep 28, 2026 BLOG ↗ RANK

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

Auth0 launches a ChatGPT plugin that provisions tenants, generates real SDK code, and queries tenant state via plain-language prompts.

NEW INTEGRATION New Auth0 plugin for ChatGPT connects to a real Auth0 tenant to create applications, configure connections, and adjust settings by plain-language prompt.

Generates tenant-matched SDK integration code instead of generic boilerplate.

Enables querying and managing tenant state — applications, users, and connections — from within a ChatGPT session.

Surfaces Auth0 documentation answers directly inside the ChatGPT conversation.

WAS THIS USEFUL? 👍 👎

Material Security

Feedback ⓘ announced · Sep 1, 2026

BLOG ↗ RANK

Material Security is an email security platform that uses AI to detect and prevent advanced threats like phishing, malware, and data exfiltration.

> Material Security significantly widened programmatic access to its detection platform, extending the Issues API and MCP server across nearly all resource types, and overhauled how analysts see why a detection fired.

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

01 Issues API and MCP server resource expansion IMPROVED

70

The Issues API now covers Issues, Messages, Accounts, Groups, Roles, OAuth Apps, and Trusted Entities as addressable resources for programmatic access. MCP server tools were broadened to work across issues, messages, accounts, and OAuth apps, including updating Trusted Entities (allowlist management) without touching the UI.

— Names all resource types but no endpoint paths or exact tool calls launch-20260901-8b2dcaca

THINNER COVERAGE BELOW

02 Detection explanation overhaul on Issue Details page IMPROVED

45

The detection explanation experience on the Issue Details page was overhauled to give analysts specific technical evidence for why a detection fired.

— Describes intent but not the actual evidence shown launch-20260901-8b2dcaca

> announced · Sep 1, 2026 BLOG ↗
>
> Material Security expands its Issues API and MCP server to cover Issues, Messages, Accounts, Groups, Roles, OAuth Apps, and Trusted Entities.
>
> Expands the Issues API to cover Issues, Messages, Accounts, Groups, Roles, OAuth Apps, and Trusted Entities as addressable resources for programmatic access.
>
> Broadens MCP server tools to work across issues, messages, accounts, and OAuth apps — including updating Trusted Entities (allowlist management) without touching the UI.
>
> Overhauled the detection explanation experience on the Issue Details page to give analysts specific technical evidence for why a detection fired.

WAS THIS USEFUL?  

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ

09/29 September 29 – Web Application and API Update (2026-09-29) BLOG ↗ RANK

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

> DNSFilter widened access to the Branding settings page for MSPs and enriched the Unblock Requests grid with category data even for non-category blocks.

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

01 **Branding settings for MSPs without custom domain** IMPROVED 55

The Branding settings page is now extended to MSPs that do not have a custom domain configured, replacing the legacy Whitelabel form; existing logos are carried over automatically.

– Names the surface and migration behavior but no exact UI path

09/29 September 29 – Web Application an…

02 **Category data in Unblock Requests grid** IMPROVED 50

The Category column in the Unblock Requests grid now shows domain category information even when a domain was blocked for a non-category reason, such as Block List, AppAware, or Security block.

– Describes exact grid and block reasons but no config or command

09/29 September 29 – Web Application an…

09/29 September 29 – Web Application and API Update (2026-09-29) BLOG ↗

LATEST STABLE

DNSFilter expands Branding settings to more MSPs and adds richer category data to Unblock Requests reviews.

Extends the Branding settings page to MSPs without a custom domain configured, replacing the legacy Whitelabel form, with existing logos carried over automatically.

Adds domain category information to the Category column in the Unblock Requests grid even when a domain was blocked for a non-category reason (Block List, AppAware, or Security block).

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ announced · Sep 24, 2026 BLOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone shipped a new `firezone` CLI for its desktop clients, adding tunnel management, Resource visibility, and a headless sign-in path for service accounts.

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

01 firezone CLI for tunnel control and unattended sign-in NEW

90

Firezone desktop clients on macOS (v1.5.21+), Windows, and Linux (v1.5.18+) now ship with a `firezone` CLI for managing tunnel state, viewing Resources, and enabling or disabling the Internet Resource. A new `firezone connect` command reads a service-account token from stdin or the `FIREZONE_TOKEN` environment variable, enabling unattended/headless tunnel connections that were not officially supported before.

Authenticate a headless server or CI runner with a service-account token without interactive prompts.

```
$ echo "$SERVICE_ACCOUNT_TOKEN" | firezone connect
# or
export FIREZONE_TOKEN="$SERVICE_ACCOUNT_TOKEN"
firezone connect
```

— Names exact command, env var, platform versions, and a runnable example.

launch-20260924-1573d61e

THINNER COVERAGE BELOW

02 macOS subcommand for system extension management NEW

35

The `firezone` CLI exposes a macOS-only subcommand for managing the system extension.

— Mentions the subcommand exists but no flags or usage detail given.

launch-20260924-1573d61e

> announced · Sep 24, 2026 BLOG ↗
>
> Firezone desktop clients (macOS, Windows, Linux) now ship with a `firezone` CLI for tunnel control and unattended service-account sign-in.

└──▷ **TRY IT**

Authenticate a headless server or CI runner with a service-account token without interactive prompts.

```
$ echo "$SERVICE_ACCOUNT_TOKEN" | firezone connect
# or
export FIREZONE_TOKEN="$SERVICE_ACCOUNT_TOKEN"
firezone connect
```

Adds `firezone connect` command that reads a service-account token from stdin or the `FIREZONE_TOKEN` environment variable, enabling unattended/headless tunnel connections that were not officially supported before.

Adds the `firezone` CLI to macOS (v1.5.21+), Windows, and Linux (v1.5.18+) desktop clients for managing tunnel state, viewing Resources, and enabling or disabling the Internet Resource.

Exposes a macOS-only `firezone` subcommand for managing the system extension.

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ announced · Sep 27, 2026

BLOG ↗ RANK

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizations.

> DigiCert extended AI Trust Manager to bring cryptographic AI Passports to agents, models, and MCP servers running inside NVIDIA's OpenShell sandbox, adding AIBOM-based signing and scanning, a signed audit trail anchored to its AI Root of Trust, and a policy-based kill switch.

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

01 **AI Passport identity for OpenShell agents** NEW 55

DigiCert AI Trust Manager now issues AI Passports for every agent, model, and MCP server running inside an NVIDIA OpenShell sandbox, giving each a cryptographic PKI-rooted identity that OpenShell policy can act on instead of a name in a config file.

– Explains mechanism (PKI-rooted identity) but no setup steps given.

launch-20260927-44f983f0

02 **Signed audit trail for OpenShell allow/deny decisions** NEW 53

Adds signed, auditable records anchored to the DigiCert AI Root of Trust for every OpenShell allow/deny decision, with confidential-compute attestation attached where applicable.

– Names the anchoring trust root and attestation but not access steps.

launch-20260927-44f983f0

03 **Policy-based kill switch for agents** NEW 50

Adds a policy-based, automated, enforceable kill switch, powered by AI Passport identity, that lets customers stop any agent for any reason across organizational boundaries.

– Describes behaviour and scope but no invocation mechanism. launch-20260927-44f983f0

04 **AIBOM signing and scanning before launch** NEW 48

Adds signing and scanning of agents, models, and MCP servers with an AI Bill of Materials (AIBOM) before a sandbox launches them.

– Names the AIBOM format but no workflow or command detail. launch-20260927-44f983f0

announced · Sep 27, 2026 BLOG ↗

DigiCert AI Trust Manager now issues AI Passports for agents running in NVIDIA OpenShell, adding cryptographic identity and a kill switch.

Adds AI Passport issuance for every agent, model, and MCP server running inside an NVIDIA OpenShell sandbox, providing cryptographic PKI-rooted identity that OpenShell policy can act on instead of a name in a config file.

Adds signing and scanning of agents, models, and MCP servers with an AI Bill of Materials (AIBOM) before a sandbox launches them.

Adds signed, auditable records anchored to the DigiCert AI Root of Trust for every OpenShell allow/deny decision, with confidential-compute attestation attached where applicable.

Adds a policy-based, automated, enforceable kill switch — powered by AI Passport identity — that lets customers stop any agent for any reason across organizational boundaries.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.49.5 (2026-09-28) BLOG ↗ LATEST

STABLE RANK

RegScale is a compliance and security governance platform that automates regulatory assessment, documentation, and continuous monitoring across organizations.

Adds an `about` command to the Axonius v2 integration to display connected instance details.

UPDATED INTEGRATION Axonius v2 integration

▾ 1 more release in this issue · 2026-09-28

> CLI 6.49.0 (2026-09-28) BLOG ↗ STABLE
>
> RegScale CLI 6.49.0 enriches Wiz STIG/CIS findings with per-criteria observations and expands Ports & Protocols coverage to VPN and cross-account paths.
>
> Wiz STIG and CIS benchmark results now record, in each control test's observations, which configuration criteria produced the pass or fail result and what Wiz found on the host.
>
> Wiz Ports and Protocols records now include ports reachable over VPN and from another cloud account.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.28.869 (2026-09-29) BLOG ↗ RANK

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

> Runlayer's Terraform ECS module gained a standalone session-search indexer and an identity-aware sandbox egress proxy for AgentCore, while the product added Claude Sonnet 5.5 and Bedrock-backed Slack monitor support, expanded AI Watch Detect's browser inventory to 32 Chromium-based AI browsers, and restricted Shadow AI Scan key creation to the MDM wizard alongside a breaking change to flow-metric collection defaults.

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

01 Identity-aware sandbox egress proxy for AgentCore NEW

93

Adds `sandbox_egress_mode` (values `direct`, `proxy`, `proxy_only`; default `direct`) to control how AgentCore sandboxes reach the internet, with `proxy` and `proxy_only` requiring `runlayer_agentcore_network_mode = VPC`; `sandbox_egress_proxy_image_uri` to override the proxy container image (must include an explicit tag or digest); `sandbox_egress_proxy_dns_servers` (default `["1.1.1.1", "8.8.8.8"]`) for upstream DNS resolution; `sandbox_egress_proxy_platform_allow_hosts` for additional allowed host patterns (`host`, `*.host`, optional `:port`); and `sandbox_egress_proxy_egress_ports` (default `[443]`) for extra TCP ports the proxy's security group may open. A `FF_ENABLE` flag `sandbox-egress-proxy`, set via `additional_backend_env_vars`, must be enabled per organization and rolled out before switching to `sandbox_egress_mode = "proxy_only"` — doing so without the flag enabled first leaves organizations without sandbox internet access for several minutes.

Enable the identity-aware egress proxy for agent sandboxes so outbound traffic is filtered by organization policy rather than allowed directly through NAT.

```hcl
sandbox_egress_mode                  = "proxy"
runlayer_agentcore_network_mode      = "VPC"
sandbox_egress_proxy_dns_servers     = ["1.1.1.1", "8.8.8.8"]
sandbox_egress_proxy_egress_ports    = [443, 8080]
sandbox_egress_proxy_platform_allow_hosts =
["*.mycompany.internal"]
```

Enable the sandbox egress proxy for all organizations before cutting over to proxy-only mode to avoid a gap in sandbox internet access.

```hcl
additional_backend_env_vars = { FF_ENABLE = "<flags already forced>,sandbox-egress-proxy" }
```

After the backend rollout completes, enforce proxy-only egress so no sandbox run bypasses the proxy.

```hcl
sandbox_egress_mode = "proxy_only"
```

— Every config key, default, and rollout sequencing named with examples. product docs

02 **Standalone session-search indexer for ECS deployments** NEW

91

Adds `enable_session_search_indexer` to deploy a standalone Kinesis consumer that extracts searchable text, entities, and findings from hook events and writes the session-search index tables in Postgres (opt-in, requires `enable_hook_events_pipeline = true`). `session_search_indexer_config` sets ECS cpu (minimum 256) and memory (minimum 1024); `session_search_index_mode` (`off`, `shadow`, `write`) controls whether the service only parses (`shadow`), writes real indexes and enables the coverage-gap alarm (`write`), or scales to 0 while preserving checkpoint state (`off`); and `session_search_indexer_extra_env` allows passthrough env vars prefixed `SESSIONS_SEARCH_` or `FF_ENABLE` / `FF_DISABLE`, such as `SESSIONS_SEARCH_FLUSH_SECONDS` and `SESSIONS_SEARCH_MAX_ROWS_PER_SECOND`.

Enable the session-search indexer in soak mode first so you can validate extraction metrics before committing real index writes.

```hcl
enable_session_search_indexer     = true
enable_hook_events_pipeline       = true
session_search_index_mode         = "shadow"
session_search_indexer_config     = {
  cpu    = 512
  memory = 2048
}
```

Promote to write mode and tune flush frequency without changing the module, using the `FF_ENABLE` flag to bypass per-org feature gates in staging.

```hcl
session_search_index_mode = "write"
session_search_indexer_extra_env = {
  SESSIONS_SEARCH_FLUSH_SECONDS = "5"
  FF_ENABLE                     = "session-search-ingest"
}
```

— Full mechanism, all config keys, and runnable HCL examples. product docs

---

03 **Standalone Kinesis consumer services for SIEM export and analytics**

NEW 80

Adds `audit_log_siem_export_kinesis_consumer_config` to run a standalone service exporting enriched SIEM objects from the audit-log Kinesis stream (`cpu >= 256` , `memory >= 512` , `desired_count >= 1`); `session_materializer_config` for a Kinesis consumer that reads the hook-events stream and performs segment compaction (`cpu >= 256` , `memory >= 3072` , `desired_count >= 0`); `cold_analyzers_consumer_config` for a consumer reading the hook-events stream (`cpu >= 256` , `memory >= 2048` , `desired_count >= 0`); and `session_siem_export_kinesis_consumer_config` for a service writing gzipped NDJSON to a managed or external S3 bucket (`cpu >= 256` , `memory >= 2048` , `desired_count >= 1`).

Ship audit-log SIEM data to a downstream system by enabling the standalone Kinesis consumer ECS service with enough capacity for production-volume audit records.

```hcl
audit_log_siem_export_kinesis_consumer_config = {
  enabled       = true
  cpu           = 512
  memory        = 1024
  desired_count = 2
}
```

— Four named services with exact resource minimums and one example. product docs

---

04 **OTLP Telemetry organization API key permission** NEW

78

Adds a new write-only OTLP Telemetry organization API key permission for OpenTelemetry ingestion (logs, metrics, traces) from AI clients such as Claude Code, Claude Cowork, and Glean; created from the Organization API keys page and must be the key's only permission — it cannot also use the AI Watch scan or hook endpoints.

Set up OTLP telemetry ingestion from Claude Code into Runlayer by creating a write-only OTLP Telemetry API key scoped to OpenTelemetry only.

> 📍 In the Runlayer console, go to Organization API keys → Create key, select the `OTLP Telemetry` permission (must be the only permission), then configure your Claude Code exporter to send logs, metrics, and traces to your Runlayer OTLP endpoint using the generated key.

— Named permission, named clients, and exact setup steps. `1.28.845 (2026-09-28)`

---

05 **Flow-metric collection now defaults on with observability export**

75

BREAKING

With observability OTLP export on and tooling export off, backend flow-metric collection now defaults on; operators must explicitly set `flow_metrics_collection_enabled = false` in their Terraform ECS inputs to opt out of the newly-automatic collection.

Opt out of automatic flow-metric collection when running obs-only OTLP export without tooling export, to avoid unexpected data ingestion.

```hcl
flow_metrics_collection_enabled = false
```

— Clear before/after default change with an exact opt-out command. `1.28.869 (2026-09-29)`

---

06 **Pi and Oh My Pi agent integration** NEW

75

Adds a Pi integration guide enabling agents built on Pi to route MCP tools through Runlayer, enforce policy on local tools like `bash`, and record activity in Sessions; adds an Oh My Pi (`omp`) tab in the connector setup guide with a one-line `/mcp add` command that signs in and saves the connector.

Connect an Oh My Pi agent to Runlayer governance in one step, routing its MCP tools through Runlayer policy enforcement.

```
$ /mcp add
```

— Names the integration, policy scope, and a runnable command. `1.28.845 (2026-09-28)`

---

07 **VPC/PrivateLink network mode for AgentCore** NEW

74

Adds `runlayer_agentcore_network_mode` Terraform input requiring `PUBLIC` or `VPC`; `VPC` mode requires `enable_privatelink = true` and either `runlayer_agentcore_consumer_vpc_id` with

`runlayer_agentcore_consumer_private_subnet_ids`, or `runlayer_agentcore_consumer_vpc_cidr_block` with matching private and public subnet CIDR lists.

— Exact input and dependency chain, no runnable example given. product docs

08 **Bedrock-backed Claude access for Slack monitors** NEW

74

Adds `BEDROCK_ENABLED` support to the agent-event-trigger consumer, granting nonstreaming `bedrock:InvokeModel` access to Claude foundation models and deployment-account Claude inference profiles for legacy platform fallback, plus `sts:AssumeRole` grants scoped to `arn:aws:iam::*:role/runlayer-bedrock-*` and `arn:aws:iam::*:role/RunlayerBedrock*` so the ECS task role can assume saved Bedrock provider roles for Slack monitors. Updating the consumer image alone does not install the new IAM permission — the Terraform module must also be upgraded and applied.

— Named env var and IAM ARN patterns with an upgrade caveat. product docs

09 **Flow-trace collection controls decoupled from export** NEW

69

Adds `FLOW_TRACE_COLLECTION_ENABLED` to enable or disable backend flow-metric generation independently of export transport (defaults to `runlayer_central_otel_enabled || runlayer_obs_otel_enabled`), and `FLOW_TRACE_ADAPTIVE_SAMPLING_ENABLED` as an independent cost lever that sheds local step-recording work even when remote export is off.

Decouple flow-trace collection from OTel export and enable adaptive sampling to reduce cost on high-throughput paths.

```hcl
FLOW_TRACE_COLLECTION_ENABLED          = true
FLOW_TRACE_ADAPTIVE_SAMPLING_ENABLED   = true
```

— Names both env vars and default logic with a config example. product docs

10 **Extra SNS topics for critical alert fan-out** NEW

62

Adds `critical_alerts_extra_sns_topic_arns` to fan out critical alarms to extra SNS topics (e.g. an internal Slack relay), capped at 3 extra topics when ToolGuard alarms are enabled (4 when `enable_monitoring` or `critical_alerts_enabled` is off).

Route critical CloudWatch alarms to an internal Slack relay SNS topic alongside the default alerting path.

```hcl
critical_alerts_extra_sns_topic_arns = [
  "arn:aws:sns:us-east-1:123456789012:slack-critical-relay"
]
```

— Named key, cap logic, and a runnable example. product docs

THINNER COVERAGE BELOW

11 **AI Watch Detect browser inventory expanded to 32 browsers**

IMPROVED

56

AI Watch Detect now inventories 32 standalone AI browsers built on Chromium — including Perplexity Comet, Dia, Opera Neon, ChatGPT Atlas, Norton Neo, BrowserOS, Fellou, Strawberry Browser, Kahana Oasis, and kweo — on macOS, Windows, and Linux, counting a browser as installed only when its own installation is found.

— Names many browsers and detection logic but no user action. 1.28.869 (2026-09-29)

12 **Shadow AI Scan keys restricted to MDM wizard** BREAKING

55

Restricts Shadow AI Scan key creation to the MDM configuration wizard only; the Organization API keys page no longer offers the Shadow AI Scan role, and the API rejects hand-created keys with it, ensuring every key links devices and scan data to a configuration card.

— Clear restriction and rejection behavior, navigation only implied.

1.28.845 (2026-09-28)

13 **AWS workload-identity option for secret stores** NEW

52

Adds `secret_stores_allow_aws_workload_identity` input to allow AWS secret stores without a role ARN to use the backend ECS task role; intended for single-org self-hosted deployments only.

— Named key and scope note, no example given. product docs

14 **Connector event triggers open to all workspaces** IMPROVED

50

Connector event triggers — such as a new Gmail label, Linear issue, or Google Drive file — are now available for agents in every workspace without requiring an access request.

— Names trigger examples and the access-gate removal, no exact path.

1.28.869 (2026-09-29)

15 **Redirect URI approval guidance** NEW 48

Adds redirect URI approval guidance — the approval screen now shows the path `Settings → General → OAuth redirect URIs` and lets members copy a shareable link for their admin.

— Exact navigation path named, thin beyond that. `1.28.845 (2026-09-28)`

16 **Faster AI Watch scans via batching and skip logic** IMPROVED 45

Speeds up AI Watch scans by batching artifact reporting in requests of up to 25 artifacts instead of one request per artifact, reducing traffic against the Runlayer deployment; also skips unchanged plugins already held by the server for the same window as skills, cutting repeat uploads on every scheduled scan tick.

— Explains mechanism and numbers but no user-facing action. `1.28.845 (2026-09-28)`

17 **Claude Sonnet 5.5 for agents** NEW 41

Agents can now run on Claude Sonnet 5.5 via the Anthropic API or AWS Bedrock, with adaptive thinking effort settings.

— Names the model and two access paths but no setup steps. `1.28.869 (2026-09-29)`

18 **Generated chat titles in agent chat list** NEW 40

Adds generated chat titles to an agent's chat list, replacing the first prompt, with rename available from the `⋯` menu.

— Describes the UI behavior and rename path, nothing deeper. `1.28.845 (2026-09-28)`

19 **Fetch deployments button for Azure Foundry provider** NEW 39

Adds a `Fetch deployments` button when adding or editing an Azure Foundry AI provider, automatically listing models served by the endpoint and populating deployment names.

— Names the UI control and its effect, thin description. `1.28.869 (2026-09-29)`

20 **Availability-zone selection documented for ECS deployment** NEW 32

Documents availability-zone selection for the main VPC, RDS, and the AgentCore consumer VPC, with the plan failing if a listed AZ is not supported in AgentCore.

— Documentation only, no config key or example. `product docs`

└──▷ **BREAKING ON UPGRADE**

! With observability OTLP export on and tooling export off, backend flow-metric collection now defaults on — operators who did not previously collect flow metrics in this configuration must explicitly set `flow_metrics_collection_enabled = false` in their Terraform ECS inputs to opt out.

! Updating the consumer image alone does not install the new `bedrock:InvokeModel` IAM permission — the Terraform module must also be upgraded and applied, or the permission will be missing after rollout.

! Setting `sandbox_egress_mode = "proxy_only"` without first enabling the `sandbox-egress-proxy` feature flag (via `additional_backend_env_vars` / `FF_ENABLE`) and waiting for the backend rollout will leave organizations without the flag without sandbox internet access for several minutes.

! The Organization API keys page no longer offers the Shadow AI Scan role, and the API rejects hand-created keys with that role; Shadow AI Scan keys must now be minted through the MDM configuration wizard.

1.28.869 (2026-09-29) BLOG ↗ LATEST STABLE

Runlayer 1.28.869 adds AI browser inventory, obs-only flow metrics opt-out, Claude Sonnet 5.5 for agents, and connector event triggers for all workspaces.

└──▷ **USE IT**

Opt out of automatic flow-metric collection when running obs-only OTLP export without tooling export, to avoid unexpected data ingestion.

```hcl
flow_metrics_collection_enabled = false
```

Adds `flow_metrics_collection_enabled = false` Terraform ECS config key to opt out of backend flow-metric collection when observability OTLP export is on and tooling export is off (collection now defaults on in that mode).

AI Watch Detect now inventories 16 standalone AI browsers built on Chromium — including Perplexity Comet, Dia, Opera Neon, ChatGPT Atlas, Norton Neo, and BrowserOS — on macOS, Windows, and Linux, counting a browser as installed only when its own installation is found.

AI Watch Detect inventories 16 more standalone AI browsers — including Fellou, Strawberry Browser, Kahana Oasis, and kweo — on macOS, Windows, and Linux.

Agents can now run on Claude Sonnet 5.5 via the Anthropic API or AWS Bedrock, with adaptive thinking effort settings.

Connector event triggers — such as a new Gmail label, Linear issue, or Google Drive file — are now available for agents in every workspace without requiring an access request.

+1 more

└──▷ **BREAKING ON UPGRADE**

! With observability OTLP export on and tooling export off, backend flow-metric collection now defaults on — operators who did not previously collect flow metrics

in this configuration must explicitly set `flow_metrics_collection_enabled = false` in their Terraform ECS inputs to opt out.

`1.28.845 (2026-09-28)` `BLOG ↗` `STABLE`

Runlayer 1.28.845 adds OTLP API key ingestion for AI clients, Pi integration, and MDM-only Shadow AI Scan key minting.

└──▷ HOW TO FIND IT

Set up OTLP telemetry ingestion from Claude Code into Runlayer by creating a write-only OTLP Telemetry API key scoped to OpenTelemetry only.

📍 In the Runlayer console, go to Organization API keys → Create key, select the `OTLP Telemetry` permission (must be the only permission), then configure your Claude Code exporter to send logs, metrics, and traces to your Runlayer OTLP endpoint using the generated key.

Connect an Oh My Pi agent to Runlayer governance in one step, routing its MCP tools through Runlayer policy enforcement.

```
$ /mcp add
```

Adds a new write-only `OTLP Telemetry` organization API key permission for OpenTelemetry ingestion (logs, metrics, traces) from AI clients such as Claude Code, Claude Cowork, and Glean; created from the Organization API keys page and must be the key's only permission — cannot use the AI Watch scan or hook endpoints.

Restricts Shadow AI Scan key creation to the MDM configuration wizard only; the Organization API keys page no longer offers the Shadow AI Scan role, and the API rejects hand-created keys with it, ensuring every key links devices and scan data to a configuration card.

Adds a Pi integration guide enabling agents built on Pi to route MCP tools through Runlayer, enforce policy on local tools like `bash`, and record activity in Sessions.

Adds an Oh My Pi (`omp`) tab in the connector setup guide with a one-line `/mcp add` command that signs in and saves the connector.

Adds generated chat titles to an agent's chat list, replacing the first prompt, with rename available from the `⋯` menu.

`+3 more`

└──▷ BREAKING ON UPGRADE

! The Organization API keys page no longer offers the Shadow AI Scan role, and the API rejects hand-created keys with that role; Shadow AI Scan keys must now be

minted through the MDM configuration wizard.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Docker Sandboxes

Feedback ⓘ v0.46.0 NOTES ↗ RANK

Docker Sandboxes isolates applications in lightweight containers, providing secure, reproducible environments for development and deployment.

> Docker Sandboxes' v0.46.0 release adds kit-arg parameterization for cloud kit mixins, auto-restart on sandbox timeout, upgraded diagnostics and kit reliability, plus a breaking change to where secret commands execute from; docs additions cover a new `sbx daemon restart` command for reloading model provider settings.

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

01 **Parameterized kit mixins via --kit-arg** NEW 90

`sbx --cloud create` gains `--kit-arg` and `--kit-args-file` flags to pass named arguments (in `name=value` form) to v2 kit mixins at sandbox creation time, or load them from a file for repeatable CI use.

Parameterize a v2 kit mixin inline when creating a cloud sandbox — useful when you need to pin a specific version of a bundled CLI tool.

```
$ sbx --cloud create --name cloud-tools claude --kit
docker.io/<NAMESPACE>/company-cli:1.0.0 --kit-arg company-
cli.version=1.2
```

Load kit arguments from a file to keep sandbox creation commands tidy and repeatable across CI pipelines.

```
$ sbx --cloud create --name cloud-tools claude --kit
docker.io/<NAMESPACE>/company-cli:1.0.0 --kit-args-file kit-
args.txt
```

Pass kit-specific configuration values when creating a cloud sandbox, so the kit can customize its behaviour at launch.

```
$ sbx --cloud create --kit my-agent-kit --kit-arg model=gpt-4o --
kit-arg max_steps=50
```

— Runnable commands with exact flag syntax shown v0.46.0

02 **Secret commands now run from a temporary directory** BREAKING

90

Secret commands configured with `sbx secret set --command`, `sbx secret set-custom --command`, or `secrets.<name>.command` in an environment file now execute from a fresh temporary directory on the host, so relative paths such as `./credential-helper` no longer resolve from the project directory or the directory where `sbx` was run. Existing environments with secret commands will be prompted to approve a one-time plan change for the working directory on the next `sbx env run`; the execution change takes effect immediately after upgrading and restarting the daemon, even before that plan is approved.

— Exact keys and migration path named, no runnable example v0.46.0

---

03 **Daemon restart for model provider changes** NEW

85

`sbx daemon restart` reloads the Docker Sandboxes daemon after setting or changing a model provider API key or adding `llama-server` to the host `PATH`, without requiring a full restart. The daemon manages a persistent model service that starts at daemon startup, survives individual sandbox or CLI exits, and lets sandboxes share a common model store and loaded models.

— Exact command with clear mechanism described product docs

---

04 **Diagnostics upload status and socket-limit reporting** IMPROVED

85

`sbx diagnose --upload` now returns a non-zero exit status on failed diagnostics uploads, enabling reliable script-level failure detection. `sbx diagnose` output also gains a `runtime_socket_limit_bytes` field reporting the socket-path length limit used by the container runtime.

Fail a CI step immediately if a diagnostics upload is unsuccessful, rather than silently succeeding.

```
$ sbx diagnose --upload || echo 'Diagnostics upload failed — check connectivity and retry'
```

— Named field and flag with a runnable example v0.46.0

---

05 **Auto-restart cloud sandboxes on timeout** NEW

75

`sbx --cloud create --on-timeout restart` makes a sandbox that hits its time limit stop and immediately restart instead of remaining stopped.

Keep a long-running cloud sandbox alive indefinitely by auto-restarting it the moment it hits its time limit.

```
$ sbx --cloud create --on-timeout restart --kit my-agent-kit
```

— Exact flag and runnable command, minimal further detail v0.46.0

06 **Kit management reliability and skills directory support**

IMPROVED 60

Kits can now install files in the agent's skills directory even when shared skills are read-only, without triggering a read-only filesystem error. Kits added through the runtime API retain their network rules and applicable agent instructions when another kit addition recreates the sandbox container. `sbx kit add` now warns if it cannot save the updated sandbox record, explaining which kit settings could be lost and the conditions that would cause the loss.

— Names sbx kit add but no example command v0.46.0

07 **WSL2 clipboard image paste fallback** IMPROVED 60

Image paste in WSL2 falls back to the Windows clipboard when Linux clipboard tools return no image, requiring `clipboard.imagePaste` to be enabled.

— Names config key, no example given v0.46.0

THINNER COVERAGE BELOW

08 **Precise error locations for sbx env** IMPROVED 55

`sbx env` reports unrecognized environment-file keys with the file, line, and column where they were declared, including across merged files.

— Names command but no example or config key v0.46.0

09 **Playwright browser downloads under balanced network policy**

IMPROVED 55

Sandboxes created with the `balanced` network policy preset can now download Playwright browser binaries from `cdn.playwright.dev` over HTTPS.

— Names preset and domain, no command shown v0.46.0

10 **Bulk secret removal revokes running credentials** IMPROVED 45

Removing secrets in bulk now revokes credentials from running sandboxes, with failed revocations retryable even after the stored secrets have been deleted.

— Describes behavior, no command or endpoint named v0.46.0

└──▷ BREAKING ON UPGRADE

! Secret commands configured with `sbx secret set --command`, `sbx secret set-custom --command`, or `secrets.<name>.command` in an environment file now execute from a fresh temporary directory on the host. Relative paths such as `./credential-helper` no longer resolve from the project directory or the directory where you ran `sbx`. Existing environments that declare secret commands will be prompted to approve a one-time plan change for the working directory on the next `sbx env run`. The execution change takes effect immediately after upgrading and restarting the daemon, even before that plan is approved.

v0.46.0 NOTES ↗ LATEST STABLE

Docker Sandboxes v0.46.0 adds restart-on-timeout for cloud sandboxes, kit-arg passthrough, and richer CLI diagnostics.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.46.0 https://github.com/docker/sbx-releases.git
# already have the repo? check out this version:
$ git checkout v0.46.0
```

└──▷ TRY IT

Keep a long-running cloud sandbox alive indefinitely by auto-restarting it the moment it hits its time limit.

```
$ sbx --cloud create --on-timeout restart --kit my-agent-kit
```

Pass kit-specific configuration values when creating a cloud sandbox, so the kit can customize its behaviour at launch.

```
$ sbx --cloud create --kit my-agent-kit --kit-arg model=gpt-4o --kit-arg max_steps=50
```

Fail a CI step immediately if a diagnostics upload is unsuccessful, rather than silently succeeding.

```
$ sbx diagnose --upload || echo 'Diagnostics upload failed — check connectivity and retry'
```

Adds `--on-timeout restart` to `sbx --cloud create` so a sandbox that hits its time limit stops and immediately restarts instead of remaining stopped.

Adds `--kit-arg` and `--kit-args-file` flags to `sbx --cloud create` for passing values to kits supplied with `--kit`.

Adds `sbx diagnose --upload` returning a non-zero exit status on failed diagnostics uploads, enabling reliable script-level failure detection.

Adds `runtime_socket_limit_bytes` field to `sbx diagnose` output, reporting the socket-path length limit used by the container runtime.

`sbx env` reports unrecognized environment-file keys with the file, line, and column where they were declared, including across merged files.

+6 more

└──▷ **BREAKING ON UPGRADE**

! Secret commands configured with `sbx secret set --command`, `sbx secret set-custom --command`, or `secrets.<name>.command` in an environment file now execute from a fresh temporary directory on the host. Relative paths such as `./credential-helper` no longer resolve from the project directory or the directory where you ran `sbx`. Existing environments that declare secret commands will be prompted to approve a one-time plan change for the working directory on the next `sbx env run`. The execution change takes effect immediately after upgrading and restarting the daemon, even before that plan is approved.

WAS THIS USEFUL?  

Straiker

Feedback ⓘ

Straiker's Technology & Channel Partnerships for AI Agent Security (2026-09-03)

BLOG ↗ RANK

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

> Straiker launched a formal Partner Collective with AWS Marketplace availability and OpenAI Select Partner status, introduced three named AI-agent-security products (Discover, Ascend, and Defend AI), and unveiled Heimdall, a cyber-hardened sandbox for evaluating and training high-risk agents under an assumed-breach model.

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

01 **Heimdall cyber-hardened AI agent sandbox** NEW 80

Heimdall is a dedicated AI agent sandbox built on an assumed-breach premise, treating every workload and reachable component as potentially compromised. It provides dedicated Apple Silicon runner hardware supporting native macOS, Linux, and Windows evaluation environments, each starting from a fresh VM, storage volume, network identity, and credentials per attempt. Egress is enforced via an independently managed network enforcement point outside the runner, applying default-deny across IPv4, IPv6, DNS, local subnets, management networks, overlays, redirects, proxies, and caches. It issues resource-scoped, time-bounded credentials per run — including GitHub App installation tokens scoped to approved repositories with a one-hour maximum expiration — via a control-plane broker that never exposes the private key to the runner or evaluation environment. It collects tamper-resistant, off-guest telemetry (host, network, credential, policy, and lifecycle evidence) following OpenTelemetry security guidance, with each record carrying run identifier, producer identity, and sequence data to detect gaps, reordering, or substitution, and performs verified teardown confirming destruction of the VM, ephemeral volumes, routes, tokens, jobs, callbacks, service state, and external artifacts before the runner can return to an admissible state.

— Rich mechanism and named controls but no direct command or config a reader can run

Heimdall: Cyber-Hardened AI Sandbox for…

THINNER COVERAGE BELOW

02 **Partner Collective ecosystem and marketplace availability**

NEW 55

Launches the Straiker Partner Collective, a formal ecosystem bringing cloud providers, resellers, systems integrators, and ISVs together to deliver AI agent security. Makes Straiker available on AWS Marketplace, enabling procurement via private offers and

committed cloud spend drawdown, and designates Straiker an OpenAI Select Partner for building, deploying, and scaling agents on OpenAI models. Also offers a co-brandable AI risk assessment powered by Discover AI and Ascend AI, plus APIs for embedding Straiker into partner-owned services and a multi-tenant deployment model for partner-delivered engagements.

— Names AWS Marketplace, OpenAI status, APIs but no mechanism or usage path

Straiker's Technology & Channel Partner…

03 **Defend AI runtime threat detection and blocking** NEW

55

Defend AI delivers runtime threat detection and blocking, reporting 98.1% threat detection accuracy against injection attempts, data exfiltration, and off-policy behavior.

— Includes a concrete accuracy metric but no deployment mechanics

Straiker's Technology & Channel Partner…

04 **Ascend AI continuous agentic red-teaming** NEW

50

Ascend AI provides continuous red-teaming of AI agents using real-world agentic attack datasets, covering prompt injection, goal hijacking, and tool misuse.

— Names attack categories but not how testing runs

Straiker's Technology & Channel Partner…

05 **Discover AI agent and MCP server inventory** NEW

35

Discover AI performs AI agent discovery and posture management, mapping every AI agent, MCP server, and agentic workflow in a customer environment.

— States what it maps but no mechanism or setup detail

Straiker's Technology & Channel Partner…

Straiker's Technology & Channel Partnerships for AI Agent Security (2026-09-03)

BLOG ↗ LATEST STABLE

Straiker launches the Partner Collective, adding AWS Marketplace availability and an OpenAI Select Partner designation for AI agent security.

Launches the Straiker Partner Collective, a formal ecosystem bringing cloud providers, resellers, systems integrators, and ISVs together to deliver AI agent security.

Makes Straiker available on AWS Marketplace, enabling procurement via private offers and committed cloud spend drawdown.

Designates Straiker as an OpenAI Select Partner, providing enterprises a validated security path for building, deploying, and scaling agents on OpenAI models.

Introduces Straiker Discover AI for AI agent discovery and posture management, mapping every AI agent, MCP server, and agentic workflow in a customer environment.

Introduces Straiker Ascend AI for continuous red-teaming of AI agents using real-world agentic attack datasets, covering prompt injection, goal hijacking, and tool misuse.

+3 more

Heimdall: Cyber-Hardened AI Sandbox for High-Risk Agents | Straiker (2026-09-02)

BLOG ↗ STABLE

Straiker launches Heimdall, a cyber-hardened AI sandbox with independent isolation, egress enforcement, tamper-resistant telemetry, and external kill-switch controls for high

Introduces Heimdall, a dedicated AI agent sandbox built on an assumed-breach premise, treating every workload and reachable component as potentially compromised.

Provides dedicated Apple Silicon runner hardware supporting native macOS, Linux, and Windows evaluation environments, each starting from a fresh VM, storage volume, network identity, and credentials per attempt.

Enforces egress via an independently managed network enforcement point outside the runner, applying default-deny across IPv4, IPv6, DNS, local subnets, management networks, overlays, redirects, proxies, and caches.

Issues resource-scoped, time-bounded credentials per run — including GitHub App installation tokens scoped to approved repositories with a one-hour maximum expiration — via a control-plane broker that never exposes the private key to the runner or evaluation environment.

Collects tamper-resistant, off-guest telemetry (host, network, credential, policy, and lifecycle evidence) following OpenTelemetry security guidance, with each record carrying run identifier, producer identity, and sequence data to detect gaps, reordering, or substitution.

+1 more

WAS THIS USEFUL? 👍 👎

Lasso Security

Feedback ⓘ announced · Sep 29, 2026 BLOG ↗

RANK

Lasso Security protects employee GenAI use, LLM applications and AI agents with runtime guardrails, data-leak prevention and CPU-only guardrail models.

> Lasso Security launched a Claude Enterprise integration built on Claude's Compliance API, bringing user activity, conversation content, and policy violations from Claude into its AI discovery and monitoring inventory.

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

01 **Claude Enterprise integration via Compliance API** NEW

80

Lasso adds support for connecting a Claude Compliance API key, pulling from the Compliance API's `activities` endpoint to ingest Claude Enterprise and Claude Platform user activity into Lasso's inventory. For Claude Enterprise organizations, it also ingests message content for each chat once a Compliance Access Key minted by the primary owner is added, enabling conversation-level inspection for sensitive data, secrets, and policy violations. Each Claude user is surfaced in Lasso's AI discovery inventory alongside existing agents and AI applications, with a per-user timeline of MCP servers connected and skills invoked; Claude policy violations are routed to a connected SIEM or SOAR platform for triage alongside alerts from other AI platforms, and a durable audit trail of org-level Claude events is stored for SOC 2, ISO, and GDPR evidence.

— Names endpoint and access-key mechanism but no exact setup command.

launch-20260929-9fec8c91

announced · Sep 29, 2026 BLOG ↗

Lasso adds a Claude Enterprise integration via Claude's Compliance API to surface user activity, MCP servers, and conversation content.

Adds support for connecting a Claude Compliance API key to Lasso, pulling from the Compliance API's `activities` endpoint to ingest Claude Enterprise and Claude Platform user activity into Lasso's inventory.

For Claude Enterprise organizations, ingests message content for each chat when a Compliance Access Key minted by the primary owner is added to Lasso, enabling conversation-level inspection for sensitive data, secrets, and policy violations.

Surfaces each Claude user in Lasso's AI discovery inventory alongside existing agents and AI applications, with a per-user timeline of MCP servers connected and skills invoked.

> Routes Claude policy violations to a connected SIEM or SOAR platform for triage alongside alerts from other AI platforms.
>
> Builds a durable audit trail of org-level Claude events and user activity, stored alongside the rest of the AI environment for SOC 2, ISO, and GDPR evidence.

WAS THIS USEFUL?  

Noma Security

Feedback ⓘ announced · Sep 28, 2026 BLOG ↗

RANK

Noma Security discovers an enterprise's AI applications and agents, manages their posture and protects them at runtime with compliance reporting.

> Noma extends its agent security platform from cloud/app environments to employee endpoints, adding endpoint discovery, a governed approval registry, identity-aware access controls, and expanded AI-DR runtime threat detection for coding and productivity agents.

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

01 Identity-aware, tool-level Access Control policies NEW

60

Adds identity-aware Access Control policies that attribute each endpoint agent to the human behind it, linking to IdP users and groups for per-user, per-group, or org-wide enforcement. Also adds tool- and action-level granularity, allowing read access to be broadly permitted while create, update, or delete operations are restricted to specific groups.

— Names IdP linkage and CRUD-level restriction scope launch-20260928-380132fb

02 AI-DR runtime threat protection for endpoint agents NEW

60

Adds AI-DR runtime threat protection for endpoint agents, detecting prompt injection (including encoded, obfuscated, and indirect injection via tool responses), sensitive data leakage (PII, PCI, secrets), malicious intent, tool poisoning, scope violations, and behavioral drift.

— Lists many named threat classes but no config/API surface launch-20260928-380132fb

THINNER COVERAGE BELOW

03 Endpoint agent discovery via EDR/MDM NEW

50

Extends Noma's agent inventory to employee endpoints, using existing EDR or MDM deployments to discover agents, MCP servers, skills, connectors, and active processes on managed devices.

— Names EDR/MDM mechanism but no exact commands or UI path launch-20260928-380132fb

04 Contextual Policies for multi-step attack correlation NEW

50

Adds Contextual Policies within AI-DR that correlate prompts, tool calls, tool responses, and reasoning across full sessions to surface multi-step attacks such as indirect prompt injection followed by secret exfiltration.

— Explains mechanism and example attack but no interface detail launch-20260928-380132fb

05 **Open Enforcement extended to endpoints** NEW 50

Extends Open Enforcement to the endpoint, enabling the same Access Control, Agent Boundaries, and AI-DR policies to be enforced through agent hooks, AI and MCP gateways, SDKs, direct APIs, and existing EDR/MDM infrastructure.

— Names enforcement surfaces but no specific config or endpoint launch-20260928-380132fb

06 **Governed registry for discovered agents** NEW 40

Adds a governed registry that automatically assigns every discovered agent, MCP server, and skill a status of approved, needs review, or blocked.

— Describes states but not how review/blocking is triggered launch-20260928-380132fb

07 **Adaptive Context Retrieval for detector tuning** NEW 40

Adds Adaptive Context Retrieval, which learns an application's traffic patterns to distinguish reusable context (system prompts, session history, RAG context) from user-controlled input, improving detector signal and reducing false positives.

— Explains purpose but no tunable settings named launch-20260928-380132fb

08 **Coverage for coding and productivity agents** NEW 40

Adds coverage for coding and productivity agents including Claude Code, Claude Cowork, Cursor, Codex, Windsurf, Kiro, Antigravity, and OpenClaw.

— Names all supported agents but no mechanism detail launch-20260928-380132fb

announced · Sep 28, 2026 BLOG ↗

Noma extends its agent control plane to employee endpoints, discovering coding and productivity agents via existing EDR/MDM infrastructure.

Extends Noma's agent inventory to employee endpoints, using existing EDR or MDM deployments to discover agents, MCP servers, skills, connectors, and active processes on managed devices.

Adds a governed registry that automatically assigns every discovered agent, MCP server, and skill a status of approved, needs review, or blocked.

Adds identity-aware Access Control policies that attribute each endpoint agent to the human behind it, linking to IdP users and groups for per-user, per-group, or org-wide enforcement.

> Adds tool- and action-level Access Control granularity, allowing read access to be broadly permitted while create, update, or delete operations are restricted to specific groups.
>
> Adds AI-DR runtime threat protection for endpoint agents, detecting prompt injection (including encoded, obfuscated, and indirect injection via tool responses), sensitive data leakage (PII, PCI, secrets), malicious intent, tool poisoning, scope violations, and behavioral drift.
>
> +4 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 29, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue